

Commissioning Status: August 2023

sPHENIX Collaboration

Abstract

sPHENIX began commissioning with beam on May 18, 2023, with the receipt of approval to operate by BHSO (the Brookhaven Site Office of DOE). This document details the progress made to date bringing the sPHENIX detectors and its associated infrastructure to a state in which physics running can begin.

Contents

1	Introduction	1
2	Commissioning Status Summary	1
3	Run and Operation Status	3
4	Magnet	6
4.1	Initial setup and test	6
4.2	Operation	6
5	Infrastructure	7
5.1	Global Interlock System	7
5.1.1	Plans for Run 24	9
5.2	Low Voltage, Bias and Control	9
5.2.1	Low Voltage Supplies and Distribution	9
5.2.2	Bias Supplies	11
5.2.3	Monitoring	11
5.2.4	Slow Control	12
6	DAQ and Triggers	13
6.1	The sPHENIX Data Acquisition System	13
6.1.1	The Core DAQ system	13
6.1.2	Overview of the Readout Electronics	15
6.1.3	The Global Level-1 and Timing system	17
6.1.4	Run Control and DAQ Operation	18
6.1.5	Data Transfers	20
6.1.6	Remaining Commissioning Tasks	21
6.2	Level-1 Trigger System	22
6.2.1	Commissioning Timeline	22
6.2.2	Remaining Commissioning Tasks	26
7	Global Detectors	27

7.1	MBD	27
7.1.1	MBD Components	27
7.1.2	MBD Commissioning Results	27
7.2	ZDC and SMD	30
7.2.1	ZDC Components	30
7.2.2	ZDC Signal Readout	30
7.2.3	Correlation between subsystems	32
7.2.4	SMD	32
7.3	sEPD	33
7.3.1	sEPD Components	33
7.3.2	sEPD Commissioning Results	33
7.3.3	Remaining Commissioning Tasks	34
8	Calorimeters	34
8.1	Electromagnetic Calorimeter	34
8.1.1	Pre-installation Commissioning	35
8.1.2	Cooling	36
8.1.3	Post-Installation Commissioning	36
8.1.4	Leakage Currents	39
8.1.5	Remaining Commissioning Tasks	39
8.2	Hadronic Calorimeter	41
8.2.1	Pre-Beam Commissioning	41
8.2.2	Data Acquisition, Sampling and Triggers	42
8.2.3	HCal Trigger	42
8.2.4	Cosmics	43
8.2.5	Tower by Tower Absolute Calibration Derivation	43
8.2.6	LED	45
8.2.7	Correlations with other subsystems	46
8.2.8	Remaining commissioning tasks	48
9	Tracking Detectors	48

9.1	MVTX	48
9.1.1	Installation	49
9.1.2	Cooling	50
9.1.3	Slow controls and monitoring	51
9.1.4	Calibration	52
9.1.5	Cosmic events	54
9.1.6	Beam backgrounds	56
9.1.7	Remaining Commissioning Tasks	59
9.2	INTT	61
9.2.1	Commissioning of Cooling and Nitrogen Systems	62
9.2.2	Commissioning of Low Voltage and High Voltage Systems	62
9.2.3	Commissioning of FELIX Readout System	63
9.2.4	Detector Performance with Beam Collisions Au + Au at 200 GeV	64
9.2.5	Detector Performance with Cosmic-rays	65
9.2.6	Summary	67
9.3	TPC	68
9.3.1	TPC Overview	68
9.3.2	Detector turn-on and commissioning	69
9.3.3	Commissioning datasets	71
9.3.4	High Voltage Issues and GEM Deformation	74
9.3.5	Remaining Commissioning Tasks	77
9.4	TPOT	78
9.4.1	TPOT components	78
9.4.2	Commissioning of TPOT High Voltage	79
9.4.3	Commissioning of the Front-End electronics cooling	80
9.4.4	Commissioning of Front-End electronics and Data Acquisition	80
9.4.5	TPOT working point	82
9.4.6	Correlation between subsystems	82
9.4.7	Summary	83
10	Tracking	83

10.1 Commissioning Track Reconstruction	85
11 Offline Computing and Analysis	85
11.1 Online Monitoring	85
11.2 Data Processing Workflow	86
A Remaining Commissioning Tasks	89

1 Introduction

The plan to commission the sPHENIX detector was outlined in the 2022 sPHENIX Beam Use Proposal [1]. sPHENIX commissioning with beam began on May 18, 2023 and ended on August 1, 2023. This document outlines the status of commissioning as of August 18, 2023.

2 Commissioning Status Summary

We have installed and operated all the detectors in the sPHENIX MIE and all the upgrade detectors with RHIC collisions. This document provides the detailed commissioning status of each detector and its related systems, explains which commissioning steps have not been completed at the time of writing, lays out a plan how to proceed until Run 2024, and predicts the readiness of the detector at the beginning of that next run.

We begin with the status of each detector and system:

- **Minimum Bias Detector (MBD):** the MBD has provided input to the trigger from the first store. MBD operation is stable. The measured z-vertex distribution is as expected with and without crossing angle of the beams. MBD commissioning is largely complete.
- **Trigger:** the trigger system has provided clean minimum bias triggers based on MBD Local-Level-1 (LL1) primitives since day 3 of the Run. A cosmic trigger based on HCal LL1 primitives has been used to confirm the HCal calibration. Since the end of beam operations, a cosmic coincidence trigger, again based on HCal LL1 primitives, has provided cosmic tracks through the center of the detector. The jet, photon/electron, and upsiion triggers still need to be commissioned.
- **Outer Hadronic Calorimeter (outer HCal):** the outer HCal is timed in; operation is stable with more than 99% of channels live; the readout is stable at kHz rate. The measured energy distribution is correlated with the other detectors (e.g., EMCal, MBD). An initial calibration at the EM scale has been obtained from cosmic muons. High statistics cosmic ray data for tower-by-tower calibrations and temperature corrections based on the LED data are needed to finalize the calibration.
- **sPHENIX superconducting solenoid (magnet):** the magnet operation is stable. Magnet commissioning is complete.
- **Inner Hadronic Calorimeter (inner HCal):** the inner HCal status is identical to that of the outer HCal. High statistics cosmic ray data for tower-by-tower calibrations and temperature corrections based on the LED data are needed to finalize the calibration.
- **Electromagnetic Calorimeter (EMCal):** the EMCal is timed in, operation is stable with > 99% of channels live. The readout is slower by a factor two than the HCal readout due to multiplexing of the fibers. The readout is also less stable than the HCal readout due to the increased probability of hardware or software related readout issues with the larger system ($8 \times$ the number of channels the combined outer and inner HCal), but continually improved through the commissioning run. The measured energy distribution is correlated with the other detectors (e.g., HCal, MBD). A di-photon invariant mass peak has been observed at

the π^0 mass. More data is needed to demonstrate the overall EMCal uniformity and energy resolution.

- **TPC Outer Tracker (TPOT):** the TPOT is timed in; the high voltage is stable; and the readout works reliably. The TPOT multiplicity is correlated with the other detectors, for instance the INTT. Cosmic tracks in the TPOT are correlated with tracks in the INTT and MVTX. Commissioning is largely complete.
- **Time Projection Chamber (TPC):** the TPC HV has been brought up to nominal operating voltage for both the Central Membrane and the GEMs with collisions and with and without the magnetic field. Tracks have been observed in both collision events and cosmic events. The diffuse laser system has been tested. However, it has been challenging to maintain continuous operation of the TPC with the GEMs at their nominal operating gain. We are investigating mechanical, electrical, and electronic problems which may contribute to this and which may require repair. Work continues without collisions to find a satisfactory operating point for the GEMs and to complete commissioning of the Diffuse and Line laser Systems. Zero suppression, digital currents, and single event upset mitigation are planned to be implemented in firmware in the coming months.
- **Intermediate Silicon Tracker (INTT):** the INTT is timed in and routinely read out in triggered mode. The measured multiplicity is correlated with the other detectors. The measured z vertex position is consistent with that of the MBD. Cosmic tracks in the INTT are consistent with tracks in the TPOT and MVTX. Streaming readout is being tested.
- **MAPS Based Vertex Detector (MVTX):** the MVTX has been operated with beam and is now routinely read out with cosmic triggers. In beam collision events, the multiplicity in the MVTX is correlated between the various layers. Cosmic tracks in the MVTX are consistent with tracks in the TPOT and INTT. However, with beam the MVTX readout is upset by background events on the time scale of seconds, i.e., on a level that affects the duty factor of operation. Different beam conditions have been explored in collaboration with C-AD, and it was found that the level of these background events depends on beam conditions like beam size influenced by stochastic cooling, collimation, and β^* . An automated procedure has been developed to recover the readout from these background events. There is an ongoing effort to improve the speed of this auto-recovery procedure in order to minimize the effect of background on the duty factor of the detector.
- **Data Acquisition (DAQ):** Both calorimeters and tracking detectors have been routinely read out with collisions. There is an ongoing effort to identify remaining hardware (in the ADCs), firmware, and software issues. Also, while the DAQ rate is in the kHz range for some detectors, in order to reach design rate for all detectors, testing and debugging of multi-event buffering and zero suppression continue development.
- **Zero Degree Calorimeter (ZDC) and Shower Max Detector (SMD):** readout of these detectors was added using spare calorimeter electronics, and an event trigger was developed in a way that preserved the scalers used to measure and monitor luminosity. Readout of all SMD channels from collisions remains to be completed, but the cabling and signal processing is complete.
- **sPHENIX Event Plane Detector (sEPD):** readout of one of twelve interface boards was tested with collisions and the remainder of the detector and electronics was installed after the end of the RHIC run.

All of the detector systems built as part of the DOE MIE (Major Item of Equipment) and the MVTX, INTT and TPOT upgrade detectors were completely installed by April 14, 2023. The readout of the ZDC and SMD detectors was subsequently added using spare calorimeter electronics, and the sEPD was installed on RHIC maintenance days. Operations during the commissioning of the detector is fully described in Chapter 3, but in short, commissioning of the sPHENIX detector with 200 GeV Au+Au collisions was carried out from May 18, 2023 until August 1, 2023, or about 11 weeks of operation. Before the beginning of commissioning with beam, all the detectors were tested to the extent possible without beam or cosmic ray triggers, except for the TPC, which received approval from BHSO for operation with gas on May 18, 2023 and the magnet, which was ramped to full field for the first time in the run on May 31, 2023 after completion of the RHIC cool-down. Operations continued until the end of the RHIC run with the failure of the 1004B valve box on August 1, 2023.

3 Run and Operation Status

The highest priority of the sPHENIX collaboration for the inaugural run of the sPHENIX detector was to commission all detector systems to be well-prepared for data taking in runs 2024 and 2025. This is well-aligned with the view expressed in the 2022 PAC report that, “The PAC endorses the continued support of the sPHENIX detector assembly, installation and commissioning, consistent with sPHENIX as the highest priority of the RHIC program.” Towards the end of the run, it was planned to have several weeks of physics data taking, the second highest priority. The commissioning of the ten brand new sPHENIX subsystems has had many successes and benefited from the hard work of many sPHENIX collaborators.

sPHENIX construction was completed at the end of April 2023. The 20 cryo-weeks planned for RHIC Run 2023 started on May 8, 2023 with the cool-down of the collider. The run was slated to end on September 25, 2023. In preparation for the upcoming RHIC Run 2023, around-the-clock sPHENIX two-person watch shifts began on April 11, 2023. This allowed us to continuously power the experimental racks and therefore the detector low voltage and bias voltage. With the power on, we started testing the individual detector subsystems in their final configuration and exercising the read-out. We also started testing the new PLC-based alarm and rack monitoring system. Two weeks into our commissioning period, we added a DAQ Operator as a third person to the shift crew, as we started to routinely read out larger and larger parts of the detector. Halfway into the commissioning, we started running full four-person shifts, i.e., shift leader, detector operator, DAQ operator, and data monitor.

With the beginning of shifts, we also started daily afternoon shift change meetings, which ran in parallel with daily morning construction meetings, until construction ended in the beginning of May. At that time, the daily shift change meeting became the single meeting in which all aspects of sPHENIX operation were discussed.

The experiment was granted approval to operate on May 18, 2023. Before that date we were not permitted to flow operating gas in the TPC or TPOT or to flow liquid nitrogen or liquid helium to the sPHENIX magnet. With the requisite approvals in hand, we began to flow operating gas into the TPC and TPOT, allowing for several volume changes before beginning the process of conditioning the GEMs in the TPC and then testing the TPC operation with high voltage

on. Concurrently, we began to cool the sPHENIX magnet, starting from room temperature and reaching its cryogenic operating temperature on May 25, 2023. Following a successful “hipot” test, the magnet was energized.

The nominal 12-week commissioning period began on May 18, 2023, though as noted above, there was a period of about a week between that date and when operating gas, HV, and magnet were fully available. This left six weeks for physics at the end of the run if everything went as planned.

The commissioning plan, first detailed more than two years ago, was written in a linear way, commissioning each system sequentially. Instead, we proceeded in a systematic way starting the commissioning of each system as soon as possible, thus working on several systems in parallel at any given time, adjusting for beam conditions. The commissioning plan required a lot of C-AD and how they operated RHIC. We had many special requests, and we are very pleased with the response we have gotten.

The first six weeks of commissioning went extremely well. This period was dedicated to timing in the minimum bias trigger detector, establishing a clean and efficient minimum bias trigger, timing in and operating the calorimeters, measuring the effect of radiation exposure on the dark current of the calorimeter silicon photomultipliers (SiPMs), and studying the effects of a crossing angle of the beams on the vertex distribution and luminosity.

The minimum bias detector was timed in the first night of collisions; a clean minimum bias trigger was established within three days of first collisions; and the calorimeters were timed in and worked as soon as the trigger was established. In the beginning, only a fraction of the calorimeters was read out. But over the course of the next several weeks, a larger and larger fraction of the calorimeters was read out, and the read-out stability and speed were gradually improved. Since the beginning of July, the entire calorimeters were reliably read out at kHz speed. We assessed the level of radiation exposure to the SiPMs and established that it was small and in line with expectations.

The motivation for the crossing angle is to reduce unwanted luminosity outside of the tracking detector acceptance, i.e., $|z| > 10$ cm, in order to reduce space charge in the TPC, while preserving most of the luminosity inside the acceptance, i.e., $|z| < 10$ cm. It was established that with a 2 mrad crossing angle between the two beams, the z-vertex distribution narrowed as expected to a sigma of about 8 cm. However, the overall luminosity (no vertex cut) was reduced by a factor 6–7, which is a factor 2–3 more reduction than anticipated based on C-AD projections. It is expected that some of the additional luminosity loss can be recovered with the commissioning of RHIC’s 56 MHz cavity. It may be necessary to reduce the crossing angle slightly to recover more of the luminosity.

The sEPD, funded by an NSF-MRI, was deliberately not installed before the run since it obstructs access to the detectors in the magnet bore. Instead, the sEPD scintillator panels, as well as electronics boxes for partial read-out, were installed in stages during maintenance days throughout the run.

The second half of the commissioning plan was dedicated to the initial operation of the tracking detectors and then to establish the high-rate operation of the data acquisition system. The tracking detectors were initially turned on ahead of schedule. The INTT and TPOT have been operating and read out for months.

The TPC HV has been brought up to nominal operating voltage for both the Central Membrane and the GEMs with collisions and with and without the magnetic field. Tracks have been observed in both collision events and cosmic events. Initial tests of the diffuse laser system have been performed. A fast spark protection system was installed in the beginning of August, which allows to operate the GEMs at the nominal operating voltage. Continuous operation of the TPC with the GEMs at their nominal operating gain has not been fully demonstrated, and we are investigating the causes and possible remediations. As a result, commissioning of GEM operation as well as of the Diffuse and Line Laser System have not been completed. Also, zero suppression and suppression of single event upsets in the front end electronics has not yet been implemented.

The MVTX read-out is impacted by beam background. We are working on firmware and software modifications to fast-reset the read-out in order to minimize the impact of these background events. We have also worked with C-AD to study the origin and reduce the frequency of the background events.

We have started to increase the read-out speed of the calorimeters by reducing the sample size. The next steps are to introduce multi-event buffering and zero suppression. Zero suppression of the TPC and TPOT data and the duty factor of the tracking read-out are challenges that we still need to address.

On August 1, 2023, shortly after noon, a magnet quench occurred and during the response of the quench protection system a short developed in the blue ring dipole and quadrupole chains, likely damaging the valve box in the 1004B building and causing a helium leak. It was estimated by C-AD that the repair would take at least five weeks. With seven weeks left in the run, it was decided on August 4, 2023 to start the controlled warm-up of RHIC, end Run 2023, and instead extend Run 2024 with the saved funds.

For sPHENIX, the end of beam operations for 2023 does not mean the end of commissioning. Instead, many of the remaining commissioning steps can be done without beam. This includes further development of the DAQ system towards design read-out speed by updating the firmware to introduce multi-event buffering and zero suppression; work on firmware and software to fast-reset the MVTX read-out; and characterization of the behavior of the GEMs of the TPC. In particular, the TPC work, which we had wished to do before beam operations started but could not due to the late approval to operate, has become easier with the absence of beam.

We have reduced the shift crews back to two-person watch shift levels since the operation of the detector is simpler with the absence of beam. We plan to keep the shift crews staffed until at least October 3, 2023 as originally planned so we can keep the power on and the gas flowing. We have established a cosmics trigger based on a coincidence of two regions of the top and bottom HCal, which provides cosmic muons through the center of the detector. The goal is to read-out the entire sPHENIX detector at design DAQ rate and establish muon tracks through all detector before shutting down the experiment in preparation of next year's run.

4 Magnet

4.1 Initial setup and test

Before the magnet cool-down started, the magnet went through initial checks and tests, including hipot (“high potential”) testing at 520 V by Carl Schultheiss and Kin Yip. The warm hipot testing initially showed some problems and then it was realized that this was because the insulating vacuum system for the magnet had not been kept under vacuum (as the magnet was warm). Once the insulating system was pumped down to vacuum, the hipot test was repeated and was successful.

The magnet cool-down started in the afternoon on May 19, 2023 and ended around 10pm on May 25, 2023. After the cool-down, a cold hipot test was done successfully without any issues.

The power supply group led by Pablo Rosas did their initial checks (for all the interlocks) and tested the system to run with a spare regulation board successfully. The latter means that we have a spare regulation board for the power supply of the sPHENIX magnet.

4.2 Operation

In early June, the magnet was ramped up to the operating current of 4596 A (with the set point on the Power Supply being 4620.3 A, a value that we have tested and obtained during the Magnet Mapping in November 2022). There were a few instances of fast-discharge triggered by the quench/power-supply interlocks after some period of running, sometimes for more than 77 hours or in one later instance for less than 5 hours. It was not clear at the beginning where the problem was and we spent some time digging into the quench detection program. During that time we kept the magnet at 150 A partly to reduce the electron cloud effect for RHIC and also to continue testing the quench detection program itself.

In the end, the problem was traced to an intermittent pin for the clock signal of the quench protection system in a ribbon cable in the quench detection system. When it was suddenly opened, it would cause the quench protection system to send out the quench signal to cause fast-discharge of the Magnet. We have also measured three bus joints (between superconducting busses and copper) and they were all found to be about 6-7 n Ω (which is good). After replacing the cable, there have been no more issues of the “quench interlock.”

Figure 1 shows the magnet and power supply currents and the temperature of the water-cooled-bus. The temperatures of the water in the water-cooled-buses have been stable (mostly between 28 to 29°C) reflecting the good performance of the chiller system in Building 1008C. The magnet was on from June 23 - August 1, 2023 when the RHIC cryo leak problem (in Building 1004B) eventually forced the Magnet to fast-discharge due to the cryogenic interlock. The short interruption in the middle around July 5, 2023 was because we needed to power down the Magnet for the sEPD installation inside the Magnet.

Overall, after the initial cable problem was found and fixed, the magnet operation has been a great success!

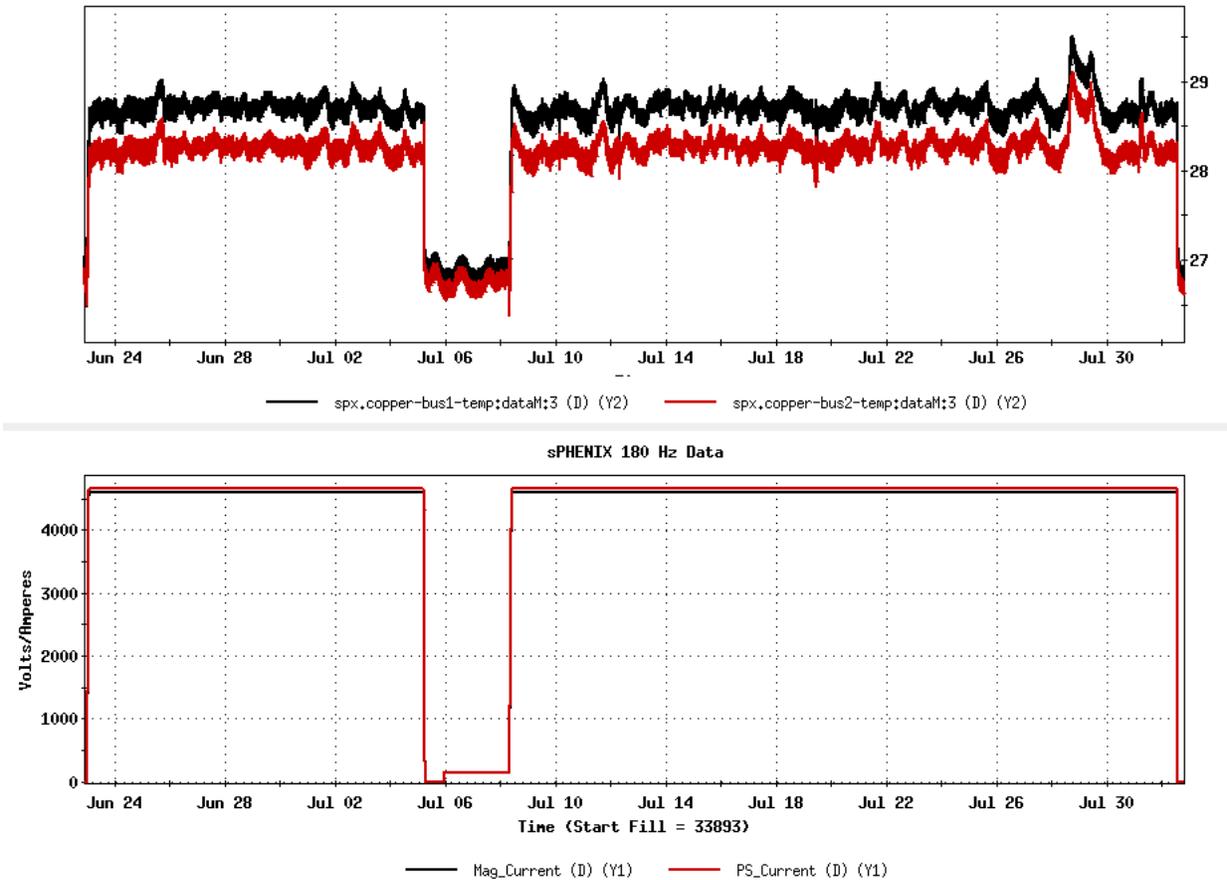


Figure 1: Top: Temperatures of water in the water-cooled-busses in °C; Bottom: Magnet and Power Supply Currents (in A) from June 23 to August 1, 2023.

5 Infrastructure

5.1 Global Interlock System

The sPHENIX Global Interlock system uses an Allen-Bradley PLC model 1769-L33ER with associated remote input/output (IO) modules to monitor and manipulate the sPHENIX detector infrastructure system. The L33ER PLC is connected to its remote IO network using a device level ring (DLR) network through a fiber optic connection. The DLR is a hardware managed network which connects to the rest of the sPHENIX slow controls network through the PLC. All DLR components are supported by a dedicated UPS and generator. If the 1008 facility were to lose power, the Global Interlock system can operate without the support of the rest of the facility's networking infrastructure. The DLR also provides a reporting mechanism to alert the control room to any communication faults within the network.

Each node of the DLR is connected to a remote IO module. There are seven nodes in total: the sPHENIX counting house, the sPHENIX control room, the service corridor in 1008A, the Gas Mixing House, the Assembly Hall, within the IR at rack 3B1, and on the mezzanine of 1008B. From these locations data is collected for the fire detection systems, HVAC for the IR and counting house,

IR environment, electronic cooling water, gas system, magnet power supply, and experiment racks.

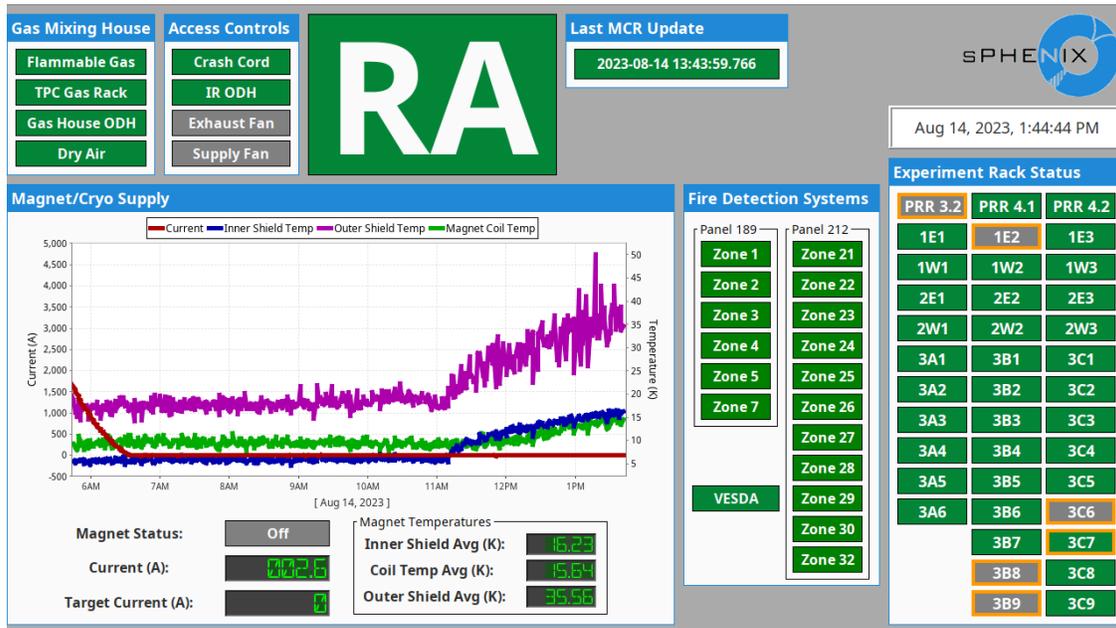


Figure 2: Annunciator screen 1 displaying fire protection, experiment rack, gas house, access controls, magnet and cryo systems.

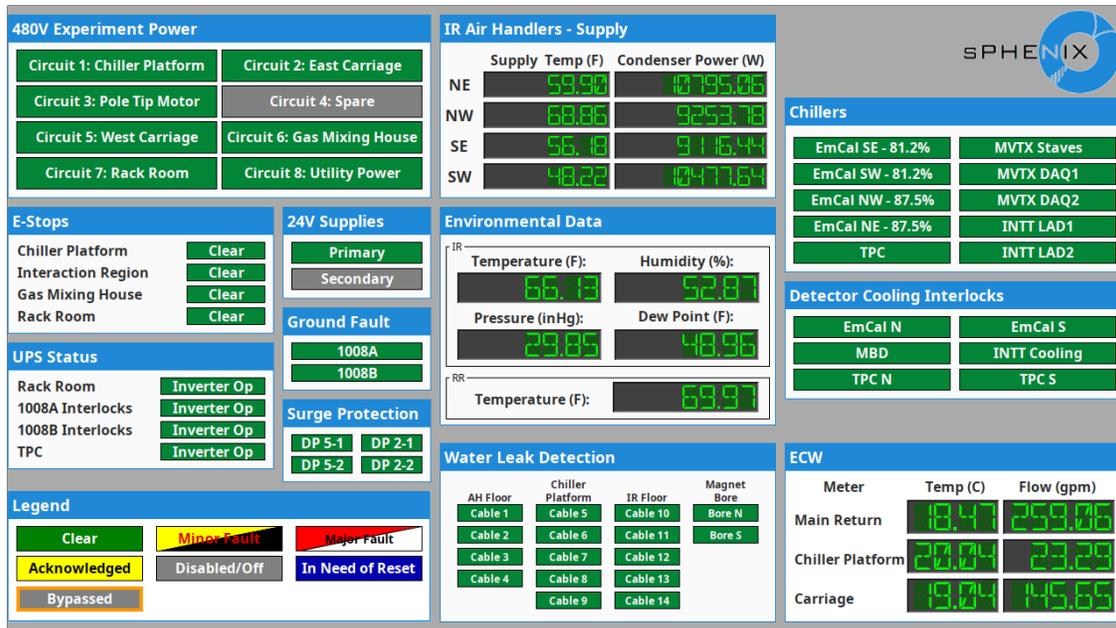


Figure 3: Annunciator screen 2 displaying experiment power, HVAC, environmental, and cooling system data.

SPHENIX infrastructure monitoring is performed by two programs, FactoryTalk View and Ignition. Ignition runs on a server within the counting house and is accessible to collaborators from outside

of 1008. It is used primarily as the control room’s alarm generator and annunciator. Two large monitors at the front of the control room display a high-level view of the systems mentioned above, with snapshots of those screens shown in Figs. 2 and 3. A one year long alarm history is kept for the Global Interlock system. FactoryTalk View is used to acknowledge and resolve any alarms that are generated by Ignition. Unlike Ignition, FactoryTalk is not accessible outside of 1008. It has writable permissions that allow operators to acknowledge, reset, bypass, and adjust alarms. A snapshot of the FactoryTalk View is shown in Fig. 4.

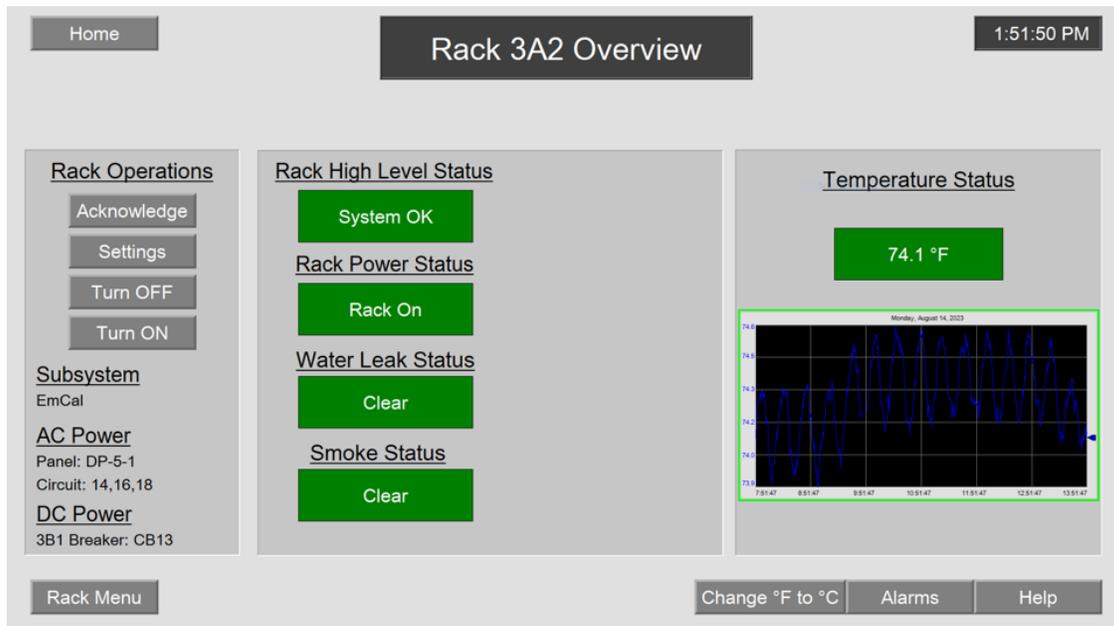


Figure 4: FactoryTalk View rack control page.

5.1.1 Plans for Run 24

The sPHENIX infrastructure team has plans to integrate a gas detection system inside the magnet bore and in the Gas Mixing House (GMH) due to the pressure and gas volume in these two areas. The installation is planned for completion prior to the 2024 run. It consists of an aspirating system connected with polytetrafluoroethylene hose lines to an infrared gas detector sensor providing monitoring of combustible gases and vapors continuously. The GMH will use an infrared sensor and will not require an aspirating system.

5.2 Low Voltage, Bias and Control

5.2.1 Low Voltage Supplies and Distribution

Low voltage for the EMCal, HCal, sEPD, TPC, MBD, Digitizers and DCM-IIs is provided from bulk power supplies from Vicor. The Mega-Pac chassis can provide up to 4 kW of power and consist of up to 8 qPAC modules that have voltage outputs ranging from 4V to 48V and are

configured for the requirements of the subsystems. Each qPAC has an enable that allows the individual qPACs to be interlocked and power cycled. Power for the EMCal, HCal sEPD and TPC is distributed using low voltage distribution system designed and built by the sPHENIX Electrical Engineering team. Each LV distribution crate consists of a controller and up to 20 LV distribution modules. The distribution modules have 8 channels, either uni-polar or bi-polar and can be individually switched and monitored for channel voltage and current through the crate controller. An example of one of the 4 crates required for the EMCal is shown in Fig. 5.

Communications with controllers is a serial protocol over ethernet and discussed below. A total of 36 Mega-Pac chassis, 296 qPACS, 15 LV distribution crates and 234 LV modules have been installed in the sPHENIX detector. Power to the Digitizer and DCM-II crates is directly wired to the qPACS using appropriate fusing. All systems were inspected and approved by a BNL EEI. Spares were procured for all components and the sPHENIX Electrical Engineering team has the ability to repair all components of the LV system. After the initial power up and testing, the LV systems have been operating reliably.



Figure 5: One of 4 low voltage bi-polar crates for the EMCal. A total of 4 crates are required for the EMCal with the HCal and sEPD each requiring an additional crate. All crates are installed in racks on the third level of the sPHENIX detector

5.2.2 Bias Supplies

The optical detectors for the EMCal, HCal and sEPD are Silicon Photomultipliers (SiPMs), that operate at a nominal 66V. Bias voltage for the SiPMs is provided by low voltage system from Weiner-ISEG. The system consists of a MPOD crate, crate controller and up to 10 LV modules with 8 channels. Each channel is capable of providing up to 125 V at 125 mA and are controlled and monitored through the crate controller. Communications with the controllers is through SNMP over ethernet. A total of 10 MPOD crates with controllers and 56 LV modules are required for the EMCal, HCal and sEPD bias systems. All components have been installed and are operational. Two of the crates installed for the EMCal are shown in Fig. 6. Sufficient spares are on hand and the sPHENIX Engineering team has a strong working relationship with the US based Weiner-ISEG electrical engineer for system support.

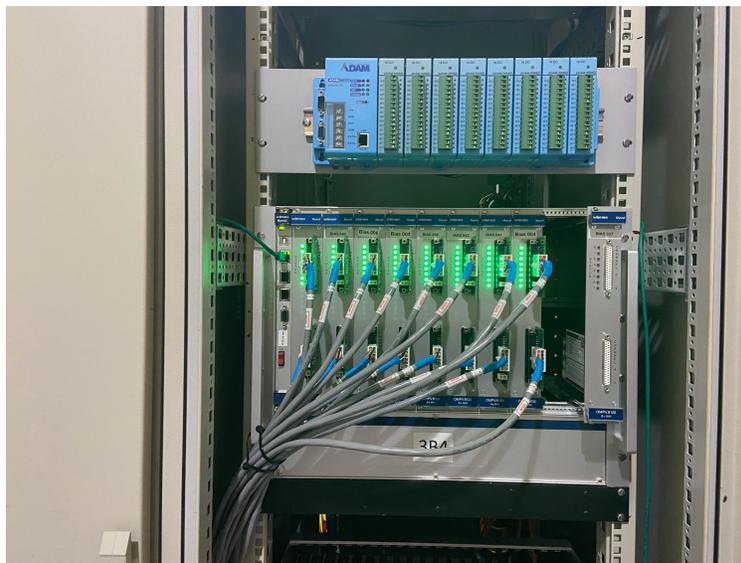


Figure 6: A crate of bias supplies for the EMCal. A total of 8 crates in 4 racks are required for the EMCal. The HCal and sEPD each require an additional crate. All crates are installed in racks on the third level of the sPHENIX detector.

5.2.3 Monitoring

There are a number of environmental variables (e.g. temperature, humidity) that need to be monitored in the different sPHENIX subsystems. For some of the subsystems monitoring electronics was integrated into the Front End Electronics. For some of the monitoring, the sPHENIX Electrical Engineering Team, designed and built custom electronics to monitor environmental variables. To monitor power supply voltages in the racks, a system based on the commercially available ADAMs system were used. These systems are readout and recorded using the sPHENIX Slow Control system.

5.2.4 Slow Control

In order to control monitor the LV and Bias systems, along with monitor the environmental conditions a slow control system has been developed. The slow control is based on commercial software and scripts developed to control the low voltage and bias systems and monitor the environment of the different subsystems. An overview of the slow control is shown in Fig. 7.

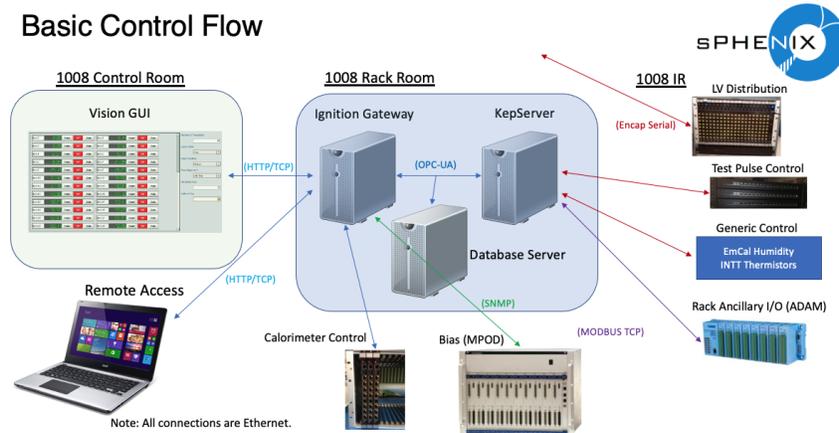


Figure 7: Pictorial representation of the slow control infrastructure for the sPHENIX detector

The Ignition Gateway serves as the central control hub and provides a graphical interface that allows for operators to control and monitor the environmental status of the sPHENIX subsystems, e.g., voltages, currents, temperatures, humidity. The Ignition Gateway also provides alarms and data logging using a Postgres database server. Communications to field devices servers and the databases uses OPC-UA protocol over ethernet.

To interface with the ADAMs systems is through a KEPSERVER that interfaces with the Ignition Gateway and communicates with the ADAMs systems using MODBUS TCP interface and Ethernet encapsulated serial communications with custom sPHENIX electronics.

6 DAQ and Triggers

6.1 The sPHENIX Data Acquisition System

The sPHENIX Data Acquisition System (DAQ) consists of several main building blocks:

- Core readout systems;
- Global Level-1 System (GL1);
- Master and Granule Timing Systems (GTM);
- Storage system;
- Networks that connect the systems.

The Global Level-1 and Timing Systems were merged into one hardware unit and are referred to as the “GL1/GTM” system.

6.1.1 The Core DAQ system

The core readout systems consist of about 60 high-end PCs that each interface to a particular component of the detector by way of a dedicated readout card. Examples of such detector components are the 24 sectors of the TPC, sectors of the inner or outer Hadronic Calorimeter, and staves of the MVTX. Each of the PCs runs an instance of the RCDAQ data acquisition system [2] that is configured individually on each of the PCs to read out the hardware in question. This is shown schematically in Fig. 8. Shown are three components from each of the main detector groups, the calorimeters (the EMCal, and the inner and outer HCal), the MBD, and the tracking detectors (MVTX, INTT, TPC, and TPOT). The PCs reading out the former group are called “Sub-Event Buffers” (SEBs), while the latter are referred to as “Event Buffer and Data Compressors” (EBDCs). Except for the different readout cards they hold, those PCs are identical and interchangeable, which simplifies the management of spares.

During the 2023 commissioning phase we used 52 readout PCs. A shortage of hardware, specifically PAR-III modules that transfer data from the DCM-IIIs to the SEB machines, resulted in two detector sections being “multiplexed”, i.e., read out with one PC. This PAR-III shortage was caused by supply chain problems procuring the parts. Additional PAR-III modules have now been received and are being tested, and thus this will be rectified in the 2024 run.

Each of the instances of the RCDAQ acquires the data from the detector component to which it is connected, and writes the data to a dedicated file. With 52 such RCDAQ instances this results in the same number of output files written concurrently, schematically shown in Fig 9. The local storage system consists of 6 so-called “Buffer Boxes”, high-end PCs with a combined number of 612 14 TB disks. Those file servers provide about 8 PB raw and 6.5 PB usable disk space with a Lustre file system.

The eventual long-term storage of the data files is done at the BNL Scientific Data and Computing Center (SDCC). Most of the data processing, analysis, and also most of the online monitoring happens there. The data are preserved indefinitely in a High-Performance Storage System (HPSS) that is tape-based and provides vast amounts of storage space.

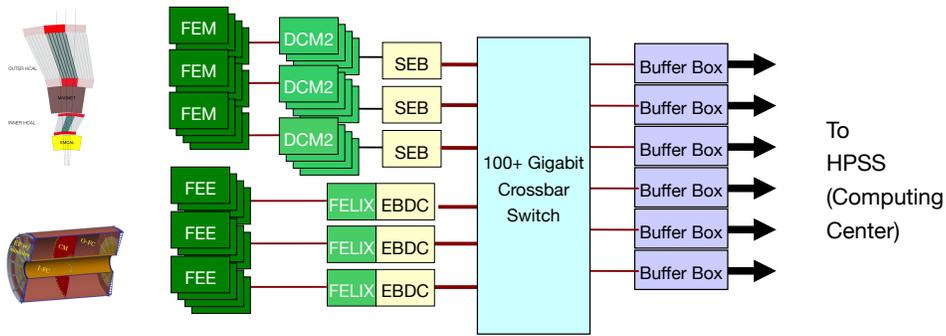


Figure 8: Schematic view of the sPHENIX DAQ system. The upper half shows the DCM-II based readout of the calorimeters and the MBD. The lower three represent the FELIX-based readout of the tracking detectors. Each component writes its files to a storage system that are connected via a high-end network switch.

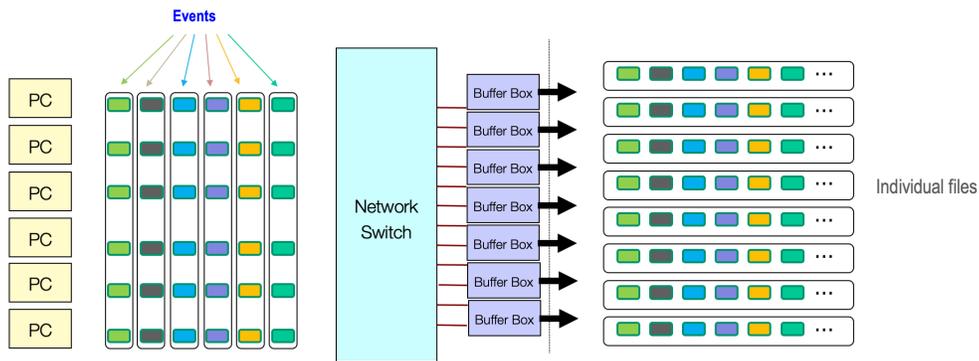


Figure 9: The principle of writing data files. Each RCDAQ component writes a series of data files to the storage system.

The analysis and most of the online monitoring is performed on a large cluster of farm nodes that access the data from a Lustre-based disk system of about 60 PB size. While it would in principle be possible to copy the data from the HPSS back to that storage system, this cannot be done on a sufficiently short time scale. It would also tie up tape drives for reading the data back that are better used for writing data to tape.

We have the network bandwidth between the experiment and the SDCC to transfer two copies of the data, one to the HPSS system, and one directly to the disk system. This enables the near-line data monitoring of the data on a timescale of about two hours, as well as a fast turnaround of our data reconstruction within 48 hours. The data transfers are shown schematically in Fig 10.

Six machines with 6 PB Lustre filesystems were commissioned and used to buffer data locally for rapid local access before transmission to SDCC. The SDCC storage systems, especially the HPSS tape storage system, work most efficiently with a steady, not greatly varying data rate. The local storage system levels (“buffers”) the varying incoming data rates and sends the steady average data rate to the SDCC. This gives the “Buffer Boxes” their name. In addition, the network between

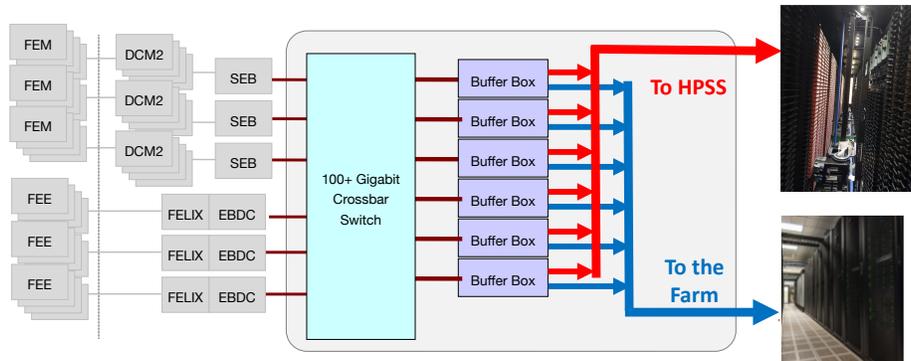


Figure 10: The data transfers from the Buffer Boxes to the HPSS storage system and SDCC disk system from where the farm nodes have access to the data.

the experiment and the SDCC does not need to be designed for the peak rates but only for the average rates, which are about 60% lower.

6.1.2 Overview of the Readout Electronics

As shown in Fig. 8, the detector systems generally fall in one of two groups. The tracking detectors (MVTX, INTT, TPC, and TPOT) are read out with the ATLAS-developed “Front-End Link eXchange” (“FELIX”) interface cards [3, 4]. The FELIX is a 16-lane PCI-Express card based on the Xilinx “Ultrascale” FPGA, and provides 48 duplex optical links with up to 10 Gbps bandwidth. A picture of the FELIX card is shown in Fig. 11. The front-end cards of the four tracking detectors that send data to the FELIX differ, and are thus described in the respective detector sections.



Figure 11: A picture of the FELIX Card (the sPHENIX production of FLX-712v2) [4]

The calorimeters, MBD, ZDC, sEPD, and the data from the Local-Level-1 systems are read out through “Data Collection Modules” (DCM-II, where the “II” refers to the 2nd generation of the system).

The ADC system for these systems is based on the Analog Devices AD9257 ADC chip with 14 bit resolution per sample. Each module shown in Fig. 12 has 4 ADC chips and 64 channels.

Digitizers for the calorimeters were tested at CU Boulder prior to shipping to BNL. Boards with errors were returned to Nevis for further testing and repair. The order only included an additional +1.6% ADC boards to account for board yield, i.e., failures, and zero spares according to Project Rules. As of December 2022, 3.7% of boards were not fully functional with many of those fixed by the start of beam. The sEPD system, funded via an NSF-MRI, did not have ADC board orders placed in time for the run. Thus, it was decided for the commissioning run period to have the high-eta south side of the EMCal not instrumented. All ADC digitizer racks are currently operational and either fully or partially populated with ADC’s.

In most systems, 3 ADC boards are read out through a custom backplane to an “XMIT” board. In Fig. 12 one can see 4 groups of 3 ADCs on the right of XMIT boards (the ones with an LED and a fiber connected). In this way, each XMIT board manages 192 calorimeter channels. Overall, the Electromagnetic Calorimeter has 32 crates, 128 XMIT boards, 384 ADC boards, and 24576 readout channels. The Inner and Outer Hadronic Calorimeter combined have 4 crates, 16 XMIT boards, 48 ADC boards, and 3072 readout channels.

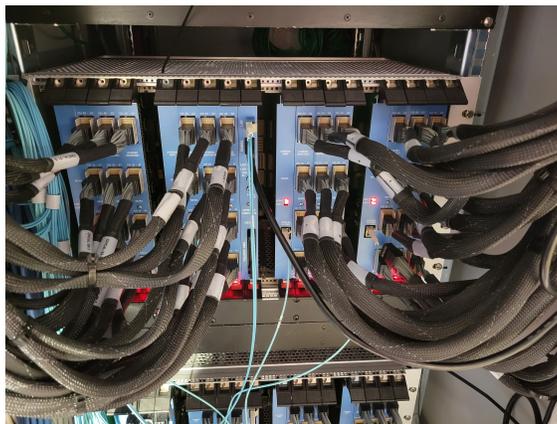


Figure 12: A Digitizer crate with the signal cables and fibers attached.

A DCM-II board supports 8 fiber inputs from the same number of XMIT boards and can thus read out 2 ADC crates. One or more DCM-II boards are then read out through a “Partitioner” board, with the data sent to the interface cards in the SEB machines. The Partitioner derives its name from the fact that it “partitions” the backplane in the crate, essentially forming independent readout units within the same crate. The individual components are configured with the Control Module seen in the leftmost slot in the DCM-II crate.

The highest bandwidth is available when just one DCM-II board is read out through a Partitioner module. As described earlier, during the commissioning phase the Electromagnetic Calorimeter was read out with two DCM-II boards per Partitioner due to a shortage of Partitioner boards.

6.1.3 The Global Level-1 and Timing system

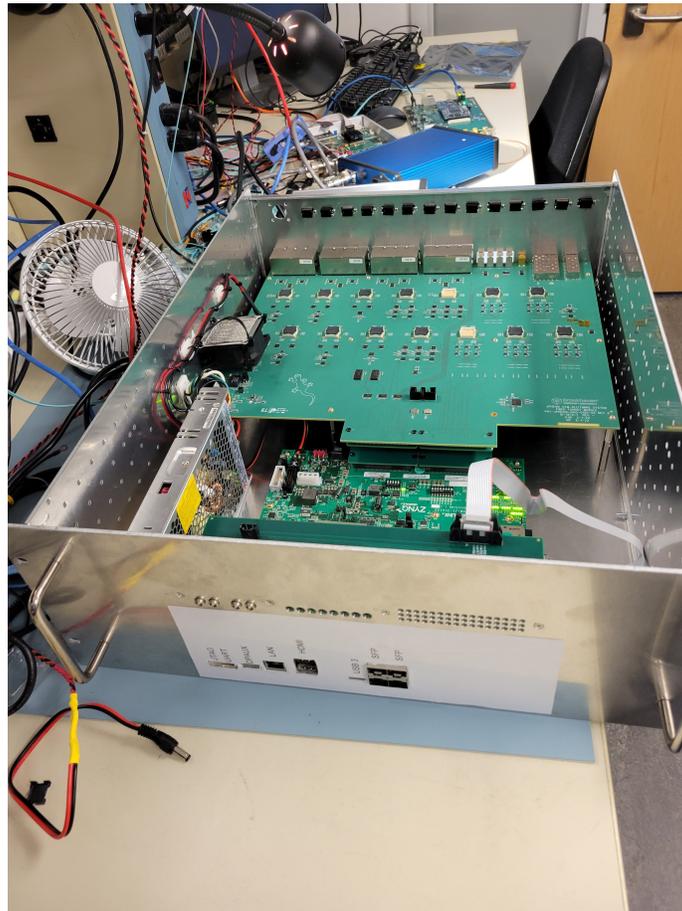


Figure 13: A picture of the GL₁/GTM unit showing the ZCU₁₀₂ FPGA board (smaller board in front) and the attached FFO board.

Originally conceived as two independent units, the Global Level-1 and the Timing Systems were merged into one hardware unit and are referred to as the “GL₁/GTM” system. The systems are based on a commercial XILINX FPGA development board, the “ZCU₁₀₂” board. It has a XILINX Zync FPGA that, in addition to the FPGA fabric, provides 2 ARM CPU hard-cores with the ability to run Linux. This architecture makes it easy to interface, configure, and interact with the system, as one can simply log in to the Linux operating system. From here one has access to the FPGA fabric with the ability to load firmware, but also to set and retrieve parameters from the FPGA, and in this way configure and control the behavior.

Fig. 13 shows one of the GL₁/GTM units partially assembled. The commercial board interfaces with a custom “Fiber-Fanout” interface board (FFO) that provides most of the experiment-specific connectivity.

Fig 14 shows the block diagram of the components of the GL₁/GTM system. The “Local-Level-1” systems shown on the left are not part of this system. They gather information from various

detectors, such as the MBD or the HCal, and provide it to the GL1 so it can make a decision whether or not to accept the data from that particular RHIC crossing for readout and storage (“trigger”). Other considerations for the GL1 to issue a trigger is the busy state of the system that determines if the data can be taken in the first place.

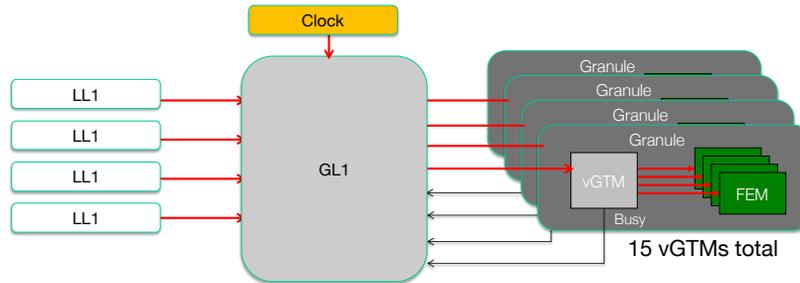


Figure 14: Block diagram of the components of the GL1/GTM system.

The GL1 receives a copy of the RHIC accelerator clock that drives the entire Collider operation, and determines when RHIC bunches cross each other and can collide. The entire GL1/GTM unit is synchronized with the RHIC clock at any time. For the Au+Au run the clock has a frequency of 56.299 MHz, which is 6 times the crossing frequency. All front-end electronics components receive this “ $\times 6$ ” clock from their vGTM and are thus also synchronized with the accelerator clock.

The GL1 interfaces with the different detector front-ends through the Timing System, which is organized as “virtual Granule Timing Modules” (vGTMs). A “Granule” refers to a detector, or a part of a detector, that has the same timing requirements *and* requires the same information for all its components to receive triggers and other assorted information. The “virtual” refers to the fact that each GTM is instantiated entirely in firmware in the GL1/GTM FPGA, not as a dedicated distinct hardware unit.

The Timing System provides 15 vGTMs that are assigned to detectors as needed. For example, both the electromagnetic and the hadronic calorimeters each need just one vGTM, while the TPC needs two, one for the south and one for the north part of the detector.

One can think of the each vGTM as an adapter of the GL1 information that fits the needs of the detector components in question, such as the kind of information provided, as well as individual timing parameters so the front-end electronics acquires the data from the selected crossing at the right moment.

6.1.4 Run Control and DAQ Operation

A server process called “Run Control” maintains connections to all DAQ components including the GL1/GTM, and orchestrates their operation.

When starting taking data (a “run”), Run Control will

- coherently configure all RCDAQ instances to take the same “kind” of data, such as physics data, calibration data, pedestals, and so on;
- issue a “begin-run” instruction to all participants so they prime themselves for data taking;
- when done, it starts the GL₁/GTM system to start acquiring data.

In order to end a run, Run Control proceeds in the reverse order:

- it stops the the GL₁/GTM system and so ends the flow of data;
- it sends all participants an end-run instruction so they flush data buffers, close files, and go into an idle state.

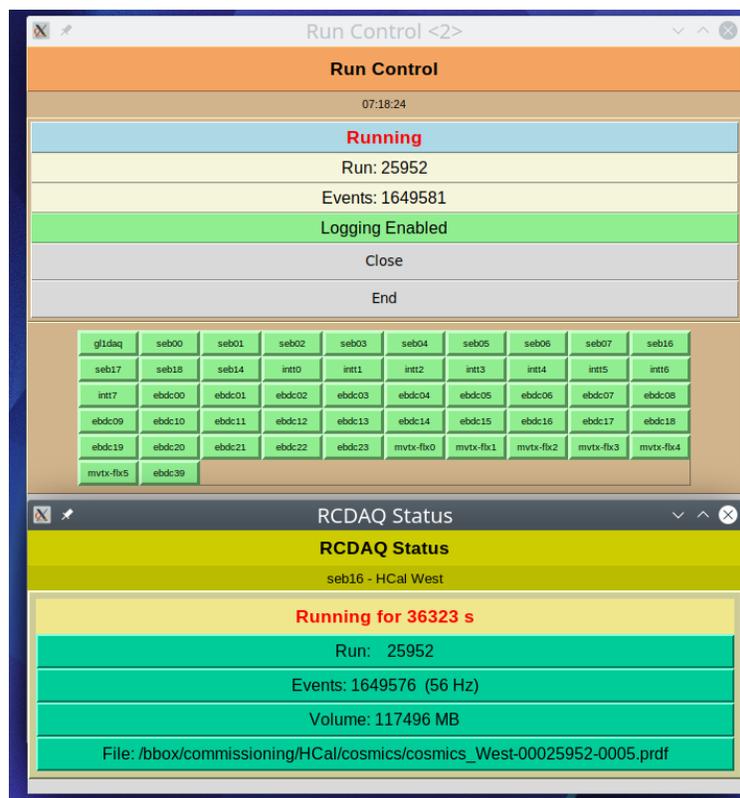


Figure 15: A screenshot of the Run Control GUI with all 52 components taking data. The lower inset is the display of one particular RCDAQ instance.

Fig. 15 shows a screenshot of the Run Control GUI with all 52 components taking data. The lower inset is the display of one particular RCDAQ instance (the west half of the hadronic calorimeter).

During the various phases Run Control logs configuration data and data about acquired event counts, run types, and so on to a database.

6.1.5 Data Transfers

The various components of the DAQ system are connected through different Ethernet networks. The network was complete and used for control and high speed data transfers throughout the run.

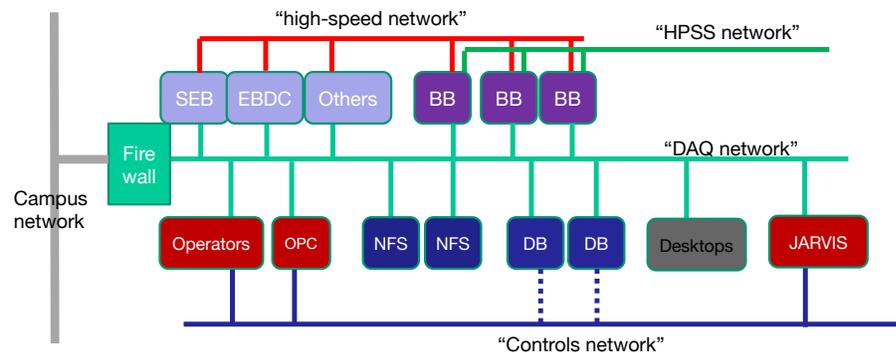


Figure 16: An overview of the different networks in the 1008 complex. All networks shown here are Ethernet networks.

Fig. 16 gives an overview of the various networks. The principal data traffic happens on the “high-speed” Network and the HPSS network. This network consists of nodes, such as the Buffer boxes, the SEBs, and the EBDCs with at least 25 Gbs network bandwidth. The SEBs and the EBDCs have 25 Gbs network interfaces, while each of the buffer boxes is connected with a 100 Gbs interface card. That gives a total theoretical bandwidth of 600 Gbs into the storage system.

The HPSS network that connects the Bldg 1008 complex with the computing center has redundant 400 Gbs bandwidth.

All components are also connected via a standard Ethernet network that is accessible from off-site (“DAQ Network”) and allows experts to intervene remotely in case of problems.

Fig. 17 shows the results of a “Mock Data Challenge” setup that was designed to tax the storage systems well beyond the envisioned data taking parameters with beam. Rather than acquiring data from the actual detector, the RCDAQ components were set up to generate data with the same properties as the detector data as fast as possible. In this way the entire front-end becomes a de-facto infinite source of data, and the performance parameters are driven solely by the storage system throughput capabilities. The left part of Fig. 17 shows the sustainable data rate first when just writing to the storage system without concurrent transfers to the SDCC. We achieved a write rate of 22.5 GByte/s this way. After starting the concurrent transfers to the SDCC, this rate dropped to 21 GByte/s. Also indicated with the green line is the bandwidth required for Heavy-Ion running. The right part of Fig. 17 shows the setup with concurrent writing and transfers over a period of 24 hours. Every hour we stopped the data taking for 3 minutes to mimic the actual operation, and also to have a visual reference to count hours in the throughput graph.

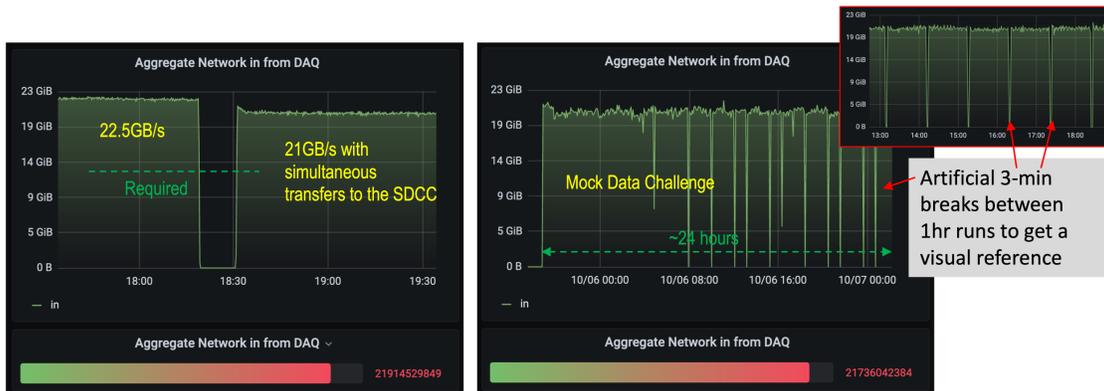


Figure 17: Left: The achieved logging rate from the DAQ to the local storage system. We first took data with no simultaneous transfers to the SDCC, and then started the transfers. We saw a sustained rate of 22.5 GBytes/s that was reduced to 21 GBytes/s in the presence of simultaneous transfers. Right: The results of a 24 hour sustained data taking and transfer, a “Mock Data Challenge”. This setup is engineered to tax the storage system well beyond the eventual data taking limits.

6.1.6 Remaining Commissioning Tasks

There are two main areas of work left to be done to achieve an overall data rate increase.

At this time, we do not zero-suppress the data, which lowers the event rate due to bandwidth limitations in the front-ends. One typically does not apply zero-suppression at the beginning of the commissioning phase and runs the data paths in a simple pass-through fashion, in order to see the raw data from the front-ends unmodified. Once one gains confidence that there are no data corruption or other problems, one begins to apply zero-suppression.

We need to apply what is known as multi-event buffering. This is a form of the general pipelining concept that allows more data to be taken by the front ends electronics before previous data have been moved to the next stage. In this way, the dead time of the DAQ can be reduced at higher trigger rate.

Most other improvements waiting to be finished fall under the general headings of improving stability, more comprehensive logging of meta-data, and more automation of the data handling, especially the transfers of data to the SDCC.

6.2 Level-1 Trigger System

The goal of the sPHENIX Trigger System is to sample the key physics (e.g., jets, photons, Y_s) from the RHIC delivered luminosity and reduce the selected event rate below the specified 15 kHz sPHENIX data acquisition bandwidth. This goal is achieved via the Local Level-1 (LL1) Trigger System providing detector inputs to the Global Level-1 (GL1) within a specified 4.2 microsecond latency (equivalent 40 ticks of the RHIC 9.4 MHz crossing clock) during which detectors are able to buffer data output either in the Front End Electronics or in the downstream Data Acquisition System.

The only systems providing such LL1 trigger primitives are the Minimum Bias Detector (MBD) and the calorimeter detectors, including the inner Hadronic Calorimeter (iHCal), outer Hadronic Calorimeter (oHCal), and Electromagnetic Calorimeter (EMCal). The Zero Degree Calorimeter (ZDC), which was not part of the original sPHENIX trigger plan, is now also included in terms of LL1 trigger primitives available.

There are details on the different physics and associated trigger requirements in Section 7.4 of the sPHENIX Technical Design Report [5]. That document details the trigger algorithms and the GEANT-4 simulation results on efficiency and rejections. In accordance with the sPHENIX Beam Use Proposal and sPHENIX Technical Design Report, for Au+Au 200 GeV collisions, the primary physics trigger is a simple interaction trigger. A coincidence of the MBD North and South with two PMTs firing on each side (with a signal above approximately 50% of the MIP) is sufficient to select $91 \pm 1\%$ of the 7.2 barn inelastic cross section - as determined from HIJING simulations run through the full sPHENIX GEANT-4. In $p+p$ and $p+Au$ 200 GeV collisions, more selective triggers utilizing calorimeter primitive inputs are required.

The LL1 design is broken into logical “design blocks” each consisting of a 570k log cell ALTERA Arria 10 Field Programmable Gate Array (FPGA), 3 Foxconn 12-port Minipod optical receiver modules, and 1 Foxconn 12-port optical transmitter module. These blocks are flexible enough to be used throughout the LL1 trigger system. The LL1 trigger system is broken into four areas: EMCal data processor, HCal data concentrator, jet trigger, and pair trigger blocks. A board layout is shown in Figure 18.

The MBD and Calorimeter ADC digitizers sample from the pulse waveform at 60 MHz. The electronics then remove baseline fluctuations and low frequency drift for every channel. The fully digital ADC output enables sending trigger primitives based on the raw ADC data to the Level-1 system. This transmission for the MBD and calorimeter systems (for every 64 channels) is via a 2.4 GHz bandwidth fiber connection. Figure 19 shows a simple schematic for the three trigger “blocks” including the basic I/O.

6.2.1 Commissioning Timeline

Ahead of Run 2023, we tested the readout chain including (i) the ADC output going to the LL1 and (ii) primitives readout from LL1 to jSEB2. This was done via test setups on the bench at both Nevis Laboratories and University of Colorado Boulder. However, these are not full system tests in the real experiment. In the real experiment, the ADC and LL1 electronics are run from different virtual Global Timing Modules (vGTM). The Level-1 output, optical or via Lemo cable

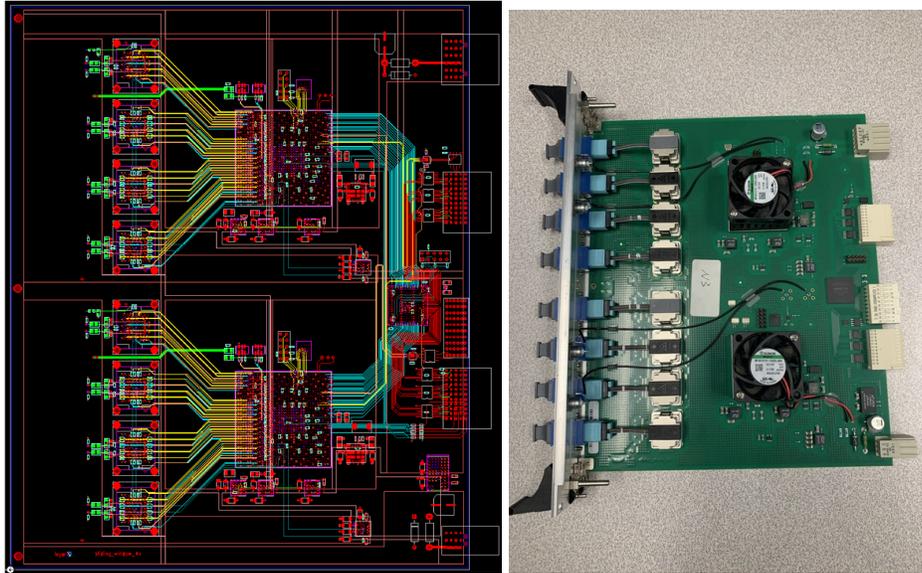


Figure 18: (left) Board layout for the sPHENIX LL₁ trigger “design block.” (right) Populated LL₁ single board.

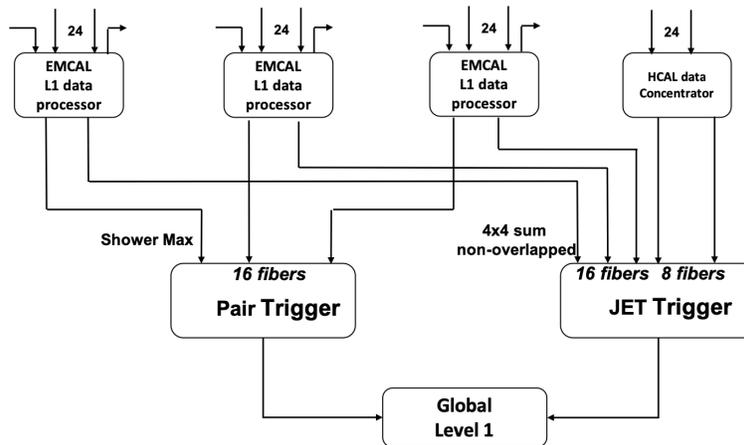


Figure 19: Simple schematic for the three LL₁ trigger “blocks.”

outputs, also needs to be received and synchronized into the GL₁.

In May 2023, communication between the MBD and LL₁ electronics was established and shortly thereafter selection based on the number of MBD channels fired on the North and South was functioning. This enabled selection on “minimum bias” Au+Au 200 GeV collisions for detector commissioning. At this point, no selection on the z-vertex of the collision was utilizing – which utilizes the additional timing information from the MBD primitives. Once the beam crossing angle of 2 mrad was established, the z-vertex Gaussian $\sigma \approx 8$ cm and no additional z-vertex selection was utilized during collision commissioning. Although the MBD only has 4 fiber inputs to the LL₁ electronics, it has all the complexities of system integration relevant to the other systems.

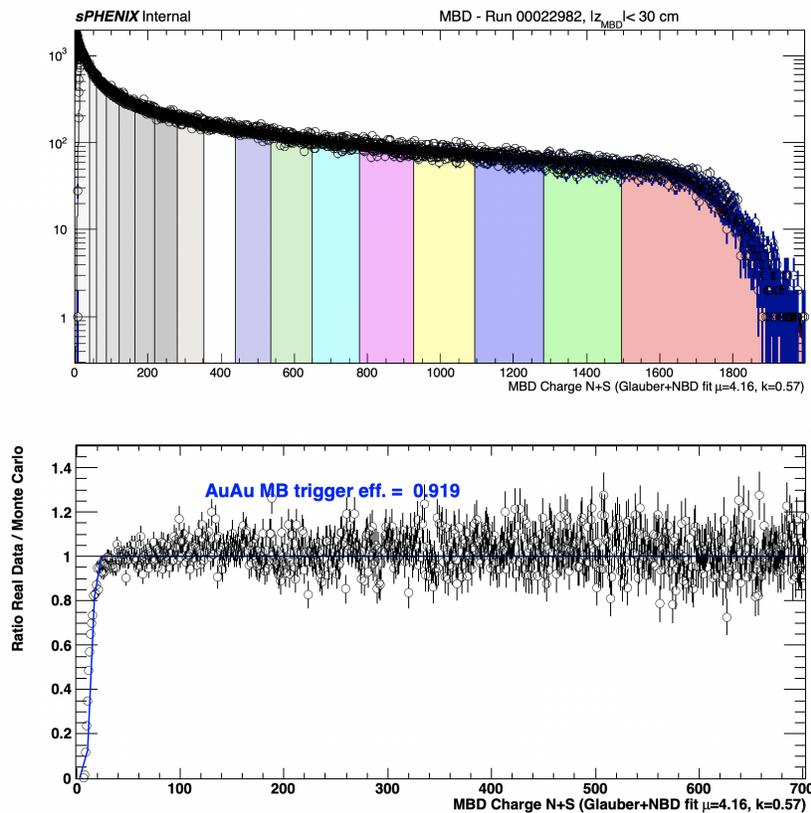


Figure 20: (Upper) MBD charge distribution corresponding to “minimum bias” triggered events requiring two PMTs to fire on each side. Also shown are the results of a Monte Carlo Glauber, folded with a negative binomial distribution response per participating nucleon. Centrality selections of 0-5%, 5-10%, 10-15%, ... 85-92% are then shown as colored regions. (Lower) The ratio of the real data to the Monte Carlo is shown. As expected, the trigger efficiency is less than 100% for events with very low MBD charge. The integral yields the 91.9% trigger efficiency.

The MBD LL1 had inputs to the GL1 initially via Lemo cable and then via fiber connection. The fiber connection provides event selection information to the GL1, which then uses a look-up table to determine if the criteria matches the user-set selection criteria. The setting of the LL1 selection digest and GL1 criteria made possible the single-side, high occupancy triggers needed for the MVTX beam background studies – see Section 9.1.6. Online monitoring for the LL1 was also available to check in real time the trigger performance.

With preliminary MBD calibrations from the sPHENIX online monitoring, the “minimum bias” selection was validated as selecting approximately 92% of the Au+Au inelastic cross section and with a trigger efficiency turn on consistent with expectations – see Figure 20.

Two different cosmic ray muon triggers were commissioned. To do that, we read out all 24 HCal trigger fibers and generated cosmic ray triggers. This is a very useful trigger for detector testing, and also important for LL1 development towards the full triggers as we use all 24 fiber inputs from HCal electronics. As such, all HCal inputs were put through the fiber patch panel and verified. One included a large OR-gate of the entire Outer HCal, where the two trigger blocks, divided in

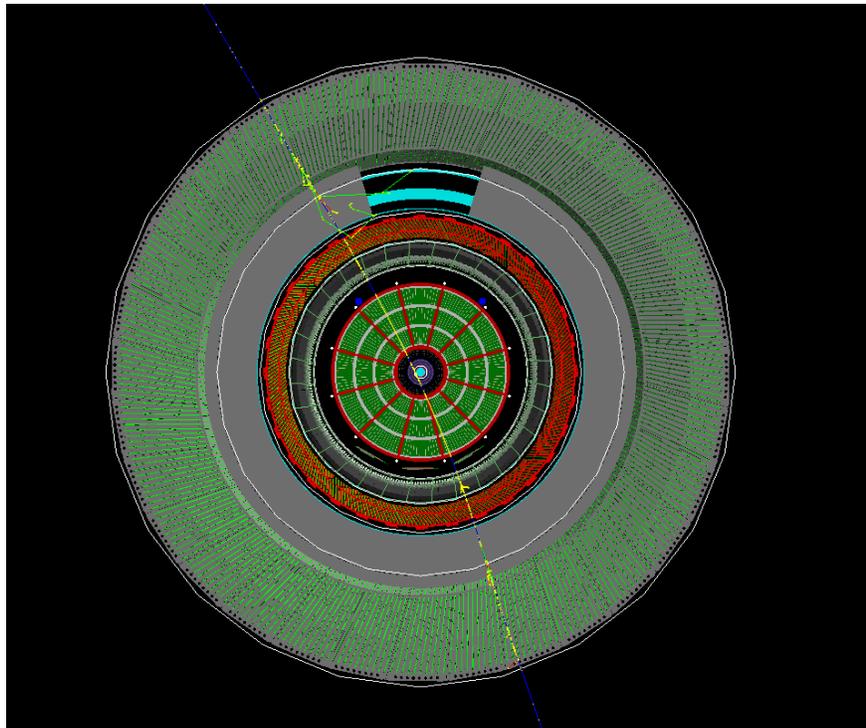


Figure 21: GEANT-4 event display from Hadronic Calorimeter cosmic trigger emulation. Shown is a muon traversing the entire suite of detectors.

sections of ϕ , are brought together in blue logic. The threshold on this trigger is was set after a series of threshold scans to optimize the signal over background. A large threshold on the HCal 2×2 sums to clear the noise threshold and to select on events where the cosmic ray muon travels through all five scintillating tiles in a HCal tower. This commissioning was successful, which enables additional calibration of the Hadronic Calorimeter – see details in Section 8 and Fig. 40.

The second cosmic ray muon trigger selected on events where the cosmic ray muon traveled through the tracking detectors. To accomplish this, the two trigger blocks were required fire in coincidence with each other. In addition to the coincidence, the areas of the Outer HCal that were not above and below the tracking detectors were masked, which further confined the regions of selection. The threshold was set to the absolute minimum of just above 6σ of the ADC pedestal, where a muon only traveling through one scintillating tile can fire the trigger. The coincidence between the two trigger blocks eliminates greatly the contamination of events without any cosmics due to noise or hot channels. This trigger was successfully commissioned and provides cosmic ray tracks through the tracking detectors – see Figure 21 from simulations.

Because of the scale, 384 input fibers from EMCAL and 24 fibers from inner HCal and 24 fibers from the outer HCal, we require a reasonable portion of front end electronics to commission and debug the full jet, photon, and pair triggers. This required the installation of trigger daughter cards, which transmit the trigger primitives, and fibers, which has been completed. The LL1 trigger system must be able to organize this large amount of data and calculate EMCAL 4×4 overlapping sums (for the photon and single electron trigger), calculate electron pair invariant masses (for the Upsilon pair trigger), and compute EMCAL/HCal combined energies in 0.8×0.8

overlapping windows in ϕ/η space (for the jet trigger). These triggers have to go through two trigger blocks; the first to concentrate input from front end inputs and the second block to deal with calculation. We highlight that all LL1 boards needed are available and all trigger cards in the ADC crates are installed.

For the first stage of processing the EMCal trigger inputs, each block received 24 fibers from the FEE system. This data block thus contains 16 channels in ϕ and 96 channels in η . That means we have complete coverage in η within this block and a partial slice in ϕ . This allows the block to calculate all 4×4 overlapping sums, except along the edges in ϕ where the data is nominally going to a different block. The electronics will send and receive the necessary cross-stitching data to the neighboring trigger processing block via additional optical transmitter and receiver ports. The overlapping 4×4 sum utilizes a sliding window algorithm. The shower maximum is determined based on a 4×4 window. The threshold and sorting used allows for us to send out the four highest shower energies per trigger block. The 4×4 non-overlapping sums are also sent out every beam crossing for use in the jet trigger. This results in one trigger output for the pair trigger module and one fiber output for the jet physics module. We will have two trigger blocks per physical board as shown in Figure 18. Thus, we will need 8 boards (16 blocks) to handle the 384 fibers incoming from the EMCal FEE.

The HCal trigger block will concentrate the HCal data from the 24 Outer HCal fibers. The 24 fiber inputs running at 2.4 Gigabits/second will have their data processed and then output the new sums in 8 fibers out running at 5.4 Gigabits/second. The incoming header words will be replaced with new headers on the sums. This concentrating is done to reduce the fiber inputs for the jet trigger block. The jet trigger block then receives this HCal data and the EMCal 4×4 non-overlapping sums to make jet patch energy sums.

The pair physics trigger block received 16 fibers from the EMCal data concentrator block. The mass calculation will be done on all combinations of electron/positron candidates using the reduced bit algorithm shown earlier.

The final trigger blocks will send input to the GL1 for final trigger decisions. In addition, the trigger accept information will be input such that trigger primitives for accepted events can be sent out to the data collection system for archiving.

6.2.2 Remaining Commissioning Tasks

Here we detail next steps that are needed for full commissioning prior to beam and with beam in 2024.

The MBD LL1 communication and primitive generation is fully functional. However, additional commissioning is needed with regards to the z-vertex selection. Most of this work relates to the MBD electronics maintaining a fixed time relative to the beam clock. In July 2023, tests started with a new FPGA code revision to pick out the maximum rather than a fixed time sample to resolve this time shift issue. No conclusive results were obtained, and so further testing even without beam, i.e., with the LASER, are needed.

The HCal LL1 inputs have all been fully tested via the cosmic ray triggers. The EMCal LL1 input fibers are in place in the patch panel, but not yet verified. Testing of the specific Jet Patch trigger

and EMCal 4x4 overlap (i.e., photon / electron) trigger algorithms remains to be done. Some work can be done with fake data without beam, and some key work will have to take place with $p+p$ collisions available. In particular loading gain constants to the ADC boards for 2x2 tower summary information sent to the LL1 and hot channel masks will require beam to generate.

The LL1 data recorded on accepted events has a data corruption problem approximately every 2500 events which occurs during the page flipping of a new direct memory access (DMA) initialization. Improvement of the plugin is in progress and can be tested without beam. The LL1 data is pushed directly to a JSEB2 board, which accepts no flow control between the LL1 board nor the JSEB2 board with the GTM, so this must be optimized in software.

At the start of Run 2024, additional focused personnel from the University of Colorado Boulder and Nevis Laboratory group will be needed for a timely completion of the outstanding items above in preparation specifically for $p+p$ 200 GeV running.

7 Global Detectors

7.1 MBD

The Minimum Bias Trigger Detector for sPHENIX primary responsibilities are to provide the data for a minimum bias trigger in heavy ion collisions to the MBD-LL1, as well as providing a z-vertex, start-time, centrality and reaction plane on an event-by-event basis. The MBD, which covers the very forward $3.5 < |\eta| < 4.5$ region, can provide a reaction plane measurement with less correlated bias than the central detectors.

7.1.1 MBD Components

The MBD re-uses the PHENIX Beam-Beam Counters (BBC), but with upgraded electronics that includes a new analog Discriminator/Shaper Front End Board that are then digitized by the same sPHENIX ADC boards used by the calorimeters. The BBC consists of two arms, each of which consists of an array of 64 Hamamatsu mesh dynode PMT's and placed at $z = \pm 250$ cm. The PMT's each have a 3 cm fused silica Cerenkov window. Those tiles are tiled to cover the radial area between 5 to 15 cm around the beam-pipe. Each channel is expected to have a time resolution of approximately 50 ps.

7.1.2 MBD Commissioning Results

From the first store provided to sPHENIX, the MBD was able to send trigger primitive data to the MBD-LL1. The MBD online monitoring was used during the commissioning to track beam conditions such as the width of the beam diamond, and could also detect when there were large amounts of beam background from upstream interactions in the triplets and other sources of beam background. The time and charge calibrations have been completed to the level required for this online monitoring, and work is on-going for final calibrations from the golden commissioning runs as a way to prepare for the upcoming physics runs in 2024 and 2025.

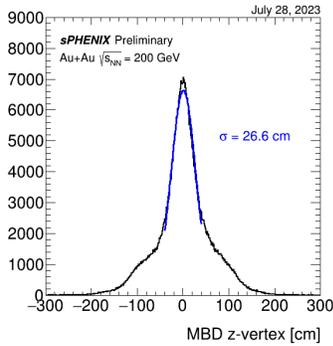


Figure 22: Z_{vtx} , Run 11947, 0 mrad crossing angle

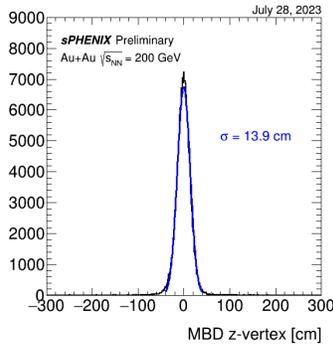


Figure 23: Z_{vtx} , Run 11949, 1 mrad crossing angle

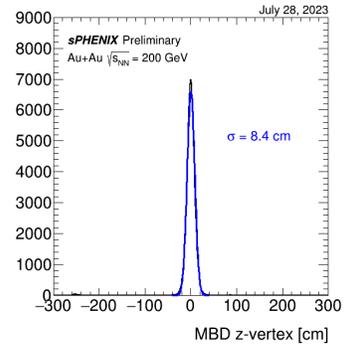


Figure 24: Z_{vtx} , Run 11950, 2 mrad crossing angle

An important priority for the MBD during the run was to verify the reduction in the z-vertex width from a crossing angle. The first runs with the crossing angle are shown in Table 1.

Run	Date	B-Field	Crossing Angle (mrad)	RMS width (cm)
11947	7/5/2023	full	0	26.6
11949	7/5/2023	full	1	13.9
11950	7/5/2023	full	2	8.4

Table 1: First crossing angle runs in sPHENIX

The reduction in width seen in Figs. 22- 24 is largely consistent with the expectations laid out in the Beam Use Proposal from 2023, and also shows that the MBD was able to provide good z-vertex measurements during the run. This has been further verified by the consistency between the z-vertex as determined by the MBD, and that from the INTT (see Fig. 60(b)).

The MBD charge calibration is done by fitting the MIP peak. The plots in Figures 25 and 26 show the distribution for the amplitude of the waveforms from PMT Channel 35 in the MBD (all other channels are similar). The prominent MIP peak comes predominantly from events where there was a single charged hadron going completely through the MBD SiO₂ Cerenkov radiator. The underlying background comes from albedo, overlaps from multiple MIPs in a PMT, or from particles which only partially go through the radiator, or have a different response than a typical hadron (such as electrons), as well as other contributions. The MIP peak is clearly seen in both field-off and field-on runs. The drop in gain of the PMT's of about a factor of 2-3 was consistent with expectations from what was known before the start of the run, and was accounted for by changes in the HV bias to the PMT's, and from being able to fit in the dynamic range of the electronics.

One unusual aspect of the sPHENIX MBD is that the albedo is different for the north and south MBD largely due to an additional flange in the direction of the north MBD. The additional flange adds material that results in conversions of photons to electron-positron pairs, which are then detected by the MBD – see Fig. 27. The original sPHENIX beam-pipe design was symmetric, but after its destruction in a warehouse fire, sPHENIX had to use an older STAR beam-pipe that required an additional flange. This leaves an asymmetry so that the north MBD sees about 30% more charge.

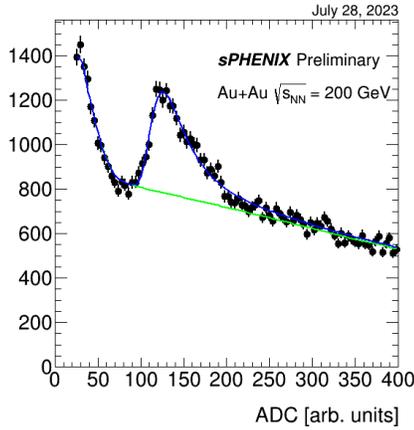


Figure 25: Run 20708, ch 35 MBD ADC Spectrum with MIP Fit

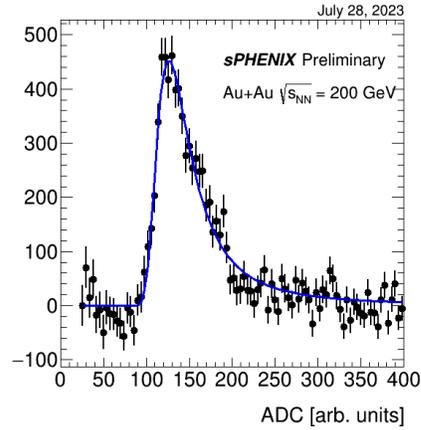


Figure 26: Run 20708, ch 35, background subtracted, Landau fit

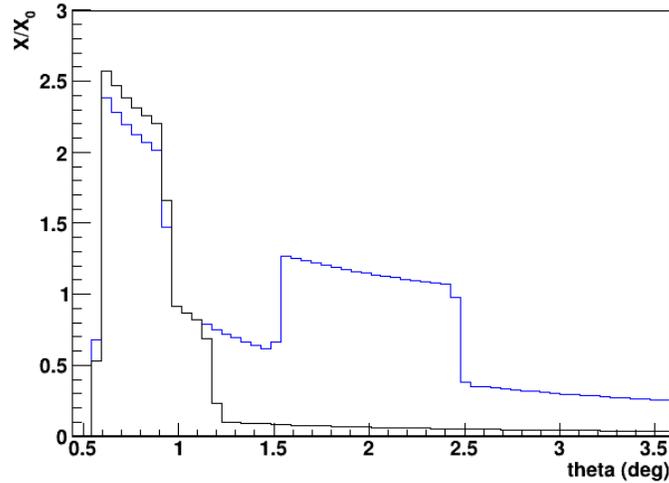


Figure 27: The integrated radiation length as a function of theta angle from the beam-pipe, for the north (blue) and south (gray) MBD. The MBD covers $1.1^\circ < \theta < 3.4^\circ$.

The asymmetry seen in real data is consistent with the sPHENIX GEANT-4 simulation, which attests well for both the quality of the charge calibration in the MBD, and for understanding the effects of the asymmetry. Figure 28 shows a radiograph from the GEANT-4 simulation with the coordinate of origin for particles leaving energy in the MBD. Studies are underway to understand if there are any implications of this asymmetry for physics, and how to best account for this north-south asymmetry in the MBD in future analyses.

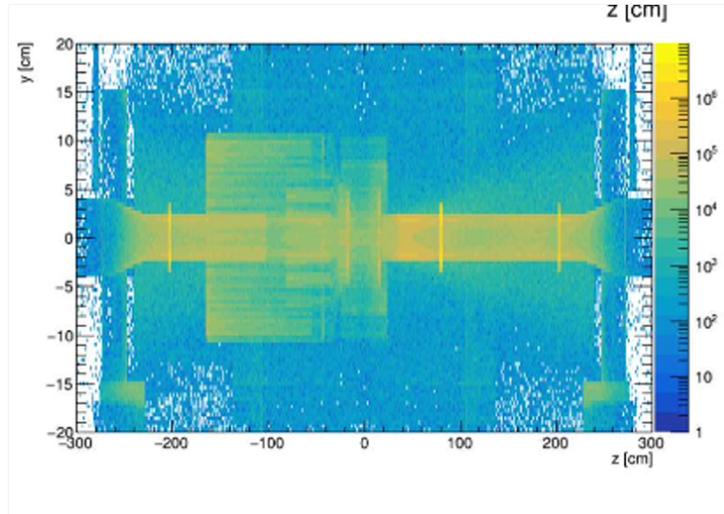


Figure 28: GEANT-4 simulation radiograph showing the coordinate of origin for particles leaving energy in the MBD.

7.2 ZDC and SMD

7.2.1 ZDC Components

While not envisioned in the sPHENIX MIE project, the sPHENIX experiment benefited from the existing PHENIX ZDC, constructed for the first RHIC running, and successfully used in PHENIX through its 2016 decommissioning. The ZDC consists of two arms, at positive and negative z , located 18m from the sPHENIX interaction point. It consists of three sections, 1.7 nuclear interaction lengths each, of 5 mm tungsten absorber plates interleaved with 0.5 mm PMMA fibers, routed to a Hamamatsu R329-2 PMT.

7.2.2 ZDC Signal Readout

The detector had not been read out since PHENIX decommissioning, but it had been utilized by the RHIC operations team to monitor collisions in the PHENIX IR. While the detector remained in place, the high voltage power supplies (for the primary ZDC PMTs) and a NIM crate (for discrimination of the individual and analog-summed signals for each side) were moved to building 1008B. The original signal cables for the ZDC to the PHENIX (now sPHENIX) electronics hall were cut, and new 130 m RG58 signal cables were run from both ZDC arms to 1008B. To integrate the system into sPHENIX, while leaving the RHIC beam monitoring undisturbed, a new analog summing board was designed by S. Boose in the sPHENIX group to provide copies of the old signals into LEMO outputs, such that they could be routed into the existing discriminator. It also provides differential outputs of each input signal over a Meritec cable output, compatible with the sPHENIX calorimeter readout. The splitter/summer board also provided individual amplifiers for the output signals and an attenuated low-gain signal path. The differential cable was routed to a new rack installed in 1008B that hosts an sPHENIX VME crate containing a single ADC and XMIT

board. New bundles of long signal and HV cables were also run from the PHENIX SMD detectors to 1008B. There, they were also converted from single-ended to differential signals using PHENIX HBD modules, and the differential signals were routed to the same ADC board as used by the ZDC. The sPHENIX ZDC was powered by the same HV values that had been used in previous PHENIX runs. No accommodation for the longer cables was made, which reduced the signal substantially, thus effectively increasing the threshold on each side. To estimate this, a special trigger signal was brought from 1008B to the sPHENIX GL₁/GTM module over optical fiber (and converted back to electrical) that could provide either a single ZDC or a coincidence. During sPHENIX commissioning, only one input could be provided, so the ZDC ran either as a single or a coincidence, but not both simultaneously.

Using the single ZDC trigger, an unbiased single neutron peak could be clearly distinguished from background (shown in Fig. 29), but with a somewhat degraded resolution from that observed in PHENIX, presumably from both radiation damage suffered over years of RHIC running, but also potentially from dispersion in the longer cables.

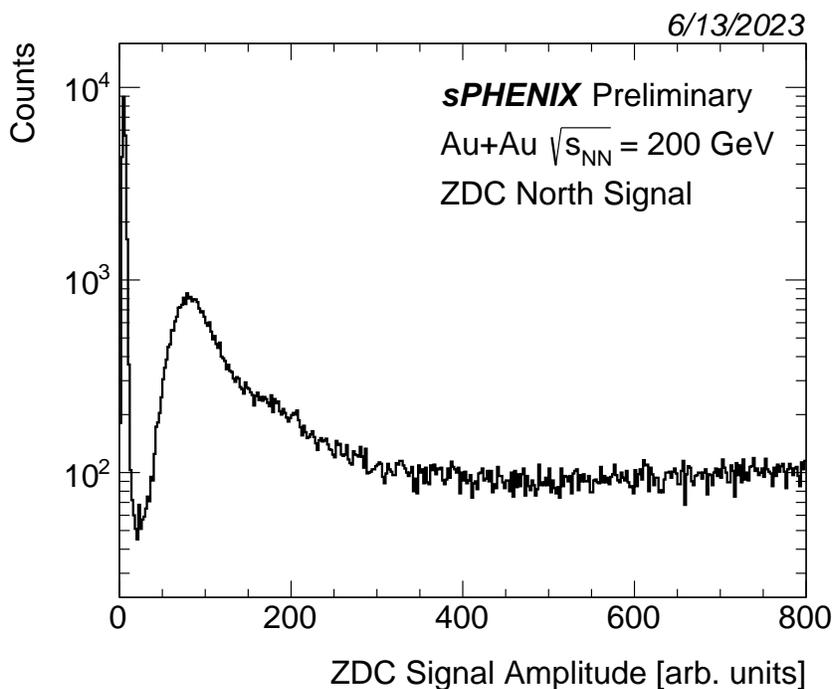


Figure 29: Energy distribution in the ZDC north arm for events triggered by the ZDC south.

Using a set of single triggers, it was estimated that the single trigger threshold (before calibration) was at about 5-6 neutrons. During a RHIC Vernier scan, this led to a measured visible cross section of about 6.5 barns (or about 60% of the 10.5b observed in previous high energy runs). Note the single trigger means on a single side (not for single neutrons), and thus the Fig. 29 data was collected by triggering on the south side (with the 5-6 neutron trigger threshold) and examining the north side.

7.2.3 Correlation between subsystems

The ZDC uncalibrated signals were sufficient to observe clear correlations of the ZDC with the MBD; as well as the ZDCs with each other (see Fig. 30), showing clear contributions from electromagnetic processes, in which one nucleus is dissociated while the other is left mostly intact, and hadronic processes (the primary one studied by sPHENIX), in which the ZDCs see the highly-correlated signals of spectator neutrons from the breakup.

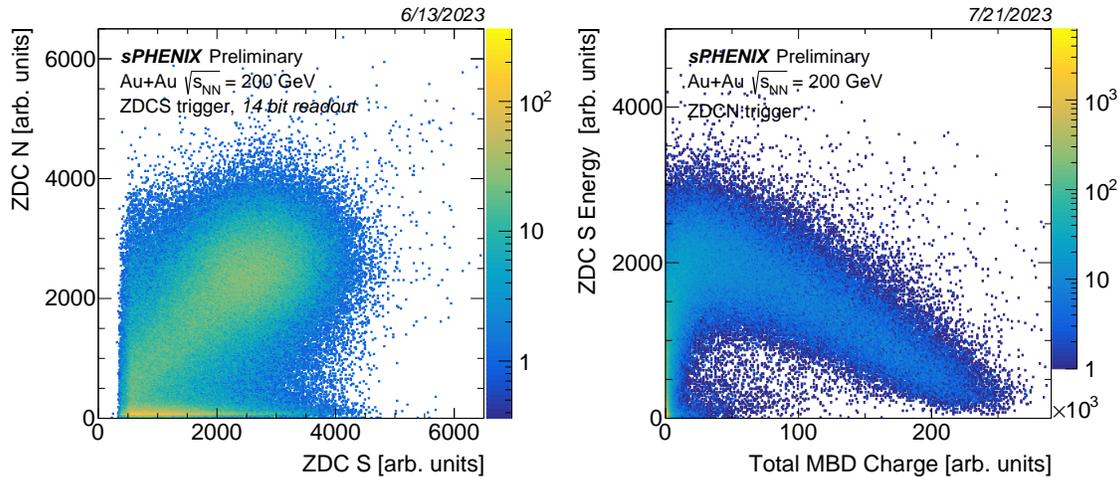


Figure 30: Left: Correlations of the energy in the south and north ZDCs for minimum bias events (electromagnetic and hadronic events). Right: ZDC south energy correlated with the total MBD charge for minimum bias events.

Calibrations are underway, utilizing both data-driven techniques to optimize the tower weights at the single neutron peak, and also techniques based on normalizing the weights to full simulations of HIJING events. The detector also appeared to be mostly stable for the duration of the sPHENIX running, but detailed time dependence studies for each detector module are underway.

7.2.4 SMD

In between the first and second section of the ZDC, a shower max detector (SMD) has been installed, with a horizontal and vertical scintillator hodoscope (8 sections vertically, 7 horizontally), read out with a multi-anode PMT. This detector was used for measuring small deflections of the measured showers. This was utilized by PHENIX for measurements of the first-order reaction plane (for flow measurements) as well as to monitor the transverse polarization of the RHIC proton beams. For use in sPHENIX, the SMD was only partly commissioned in mid-July, where half of the north SMD was digitized in the sPHENIX DAQ. The full system was readied in early August, but just after the RHIC magnet quench incident. Hence, it has not been utilized in active data taking.

7.3 sEPD

The sEPD is comprised of two scintillator disks, 1.2 cm thick, located a distance of 319 cm from the interaction region, for a total acceptance of $2.0 < |\eta| < 4.9$. Each disk has 12 sectors with 31 optically isolated tiles, for a total of 744 tiles. The ϕ segmentation is 24, except for the innermost ring which only contains 12 segments, and the η segmentation is 16. This detector will allow a forward measurement of the centrality and the event plane. The active part of the detector is essentially a copy of the STAR EPD, but the read-out electronics are the same as the sPHENIX EMCAL.

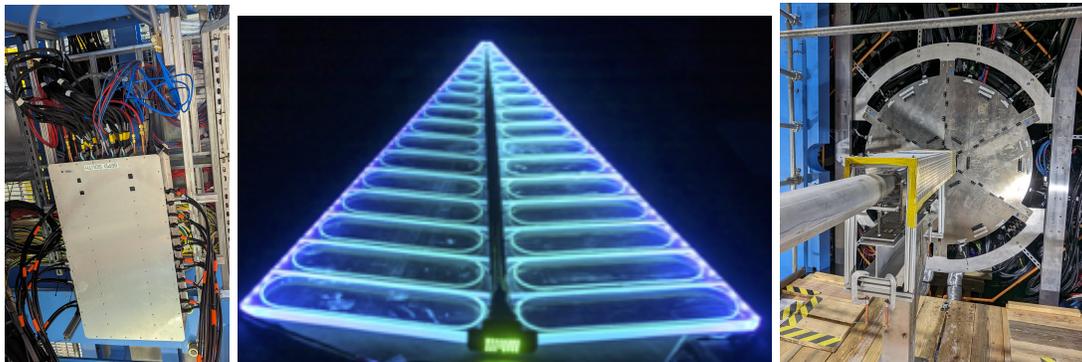


Figure 31: Photos of the read-out electronics box (left), the sEPD panel prior to light-tight wrapping (middle), and the set of sectors installed around the beam pipe (right).

7.3.1 sEPD Components

In addition to the two scintillator disks, the sEPD components also include the read-out electronics and the optical fiber connections. The signal from the sector is transported by 6.8 m long optical fibers into the sEPD SiPM boxes where they are read out by S13360-1325PE Hamamatsu SiPMs. The optical fibers split into two (15 and 16 tiles), in order to read-out in η strips. The SiPM boards have 16 channels each, so each 31 tile sector is read out by 32 channels, with the extra channel helping understand noise and verify mapping. The electronics from the SiPMs to the sub-event buffer (SEB) are common with the Calorimeter electronics. The scintillator disks and optical components were installed by July 5, 2023. A single interface board and its associated electronics were installed July 19, 2023, which allowed two sectors (in position 0 and position 5) on the north side to be read out for a total of 62 tiles. The full electronics for the South side were installed August 2, 2023, after the end of beam.

7.3.2 sEPD Commissioning Results

The sectors went through two different tests prior to the installation in the IR. The first was to use a Strontium source in order to measure the level of cross-talk and the relative strength of the signal between tiles. The second test was to use cosmics in order to confirm that we could see the MIP peak in every tile in the system installed in the IR.

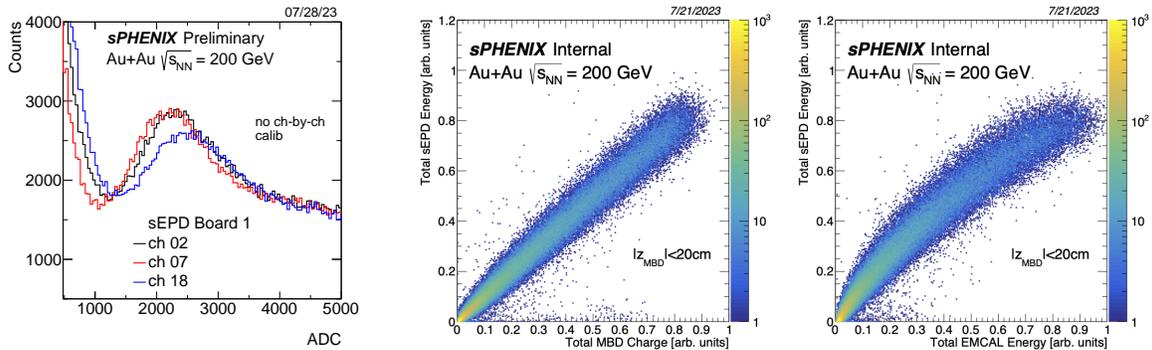


Figure 32: Left plot shows the mip peak in 3 different tiles in the high gain operation from Au+Au running. No gain matching was applied. The middle plot shows the correlation between the sum of sEPD ADCs and total charge in the MBD, normalized so that the max value was fixed to zero. The same comparison with the EMCal is shown on the right. The middle and right plots were from the normal gain operation.

After the installation of the disk, we installed a prototype SiPM Box, that allowed us to instrument two sectors on the North side. With the prototype, we were able to fully read out 61 tiles. MIP peaks for 3 tiles are shown in Figure 32. There are two modes of operation, normal gain and high gain. The mip peak plot shows the high gain operation, which did saturate the ADCs. Switching to the normal gain operation removes the saturation issue so that the correlation between the sEPD and the MBD or the EMCal, shown in 32, is linear.

7.3.3 Remaining Commissioning Tasks

Only a small fraction of the detector has been commissioned with collisions, however since everything else is a copy of what has been already installed, this should be straight forward. Some of commissioning goals can be accomplished with cosmics or with the pulsar. One key task is to confirm the mapping of the detector to ADC's, which can be by using the pulsar to set the signal and then confirm the signal shows in the channel that is expected. Sending a light signal through the fiber optics will close the loop and confirm that the mapping from the scintillator through the SiPMs to the ADCs is correct. We can also begin the calibration of the detector using cosmics to look at the mip peak over many cosmic events.

8 Calorimeters

8.1 Electromagnetic Calorimeter

The electromagnetic calorimeter (EMCal) in sPHENIX consists of a single layer just inside the inner hadronic calorimeter. The EMCal is composed of 6144 tungsten-epoxy individual modules embedded with scintillating fibers. Each module, or block, is further subdivided into four towers, the light output of which is routed to a light guide which is readout by four Silicon Photo-multipliers (SiPM). The blocks themselves are projective, meaning they tilt towards the nominal

$z = 0$ collision vertex with a slight offset to prevent channeling. The EMCal is subdivided into 64 sectors (32 on the north and 32 on the south sides of the detector). Each sector contains 6 Interface Boards (IBs), which themselves contain 16 blocks or 64 channels, for a total 24,576 readout channels. Fig. 33 shows a section drawing of an EMCal sector, including blocks, electronics PCBs, cooling, and some of the cabling. <https://www.overleaf.com/project/64b2a729a068cfd379122027>

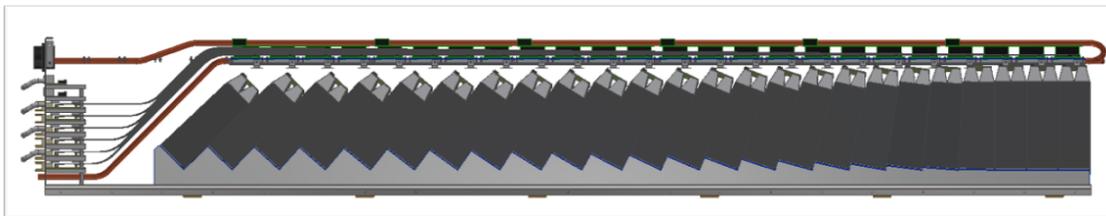


Figure 33: EMCal sector.



Figure 34: Left: An EMCal module which comprises four towers. The four light guides to the left show the delineations between the towers. Right: Fully assembled EMCal sector, including modules, SiPMs, readout electronics and cooling, with inner radius cover removed.

8.1.1 Pre-installation Commissioning

Prior to final installation in the experimental hall, a variety of steps were taken to ensure the functionality of every calorimeter module.

During fabrication, each calorimeter block was tested to ensure the desired performance. Blocks were tested to confirm that the dimensions and density were within specification to ensure maximal response uniformity when assembled into sectors. Additionally, in order to confirm that each block would be fully functional, they were tested to confirm the number of live fibers as well as the scintillation light output. As a result of this commissioning, 100% of blocks used in the final sector assembly were confirmed to be fully functional.

Fully-assembled sectors underwent a series of commissioning procedures to ensure their functionality prior to installation in the detector. These procedures included a two-week long burn-in period where the internal electronics were put under voltage with full cooling to identify any failure modes due to faulty electronic components. During this process, a small number of faults were identified and repaired. Fig. 34 shows an assembled detector module (left) and a fully assembled sector with the cover removed (right). Additionally, the physics functionality of each sector was evaluated using a combination of LED, cosmic, and single pixel studies. During this process, a small number of sectors were identified as having sub-optimal performance. These sectors were all fully repaired. At the end of this process, it was confirmed that every tower in every sector was fully functional.

8.1.2 Cooling

The EMCal uses facility chilled water (ECW) which runs at 18°C , to remove the heat generated by the electronics inside the sector enclosures. Each of the 64 sectors produces approximately 180 Watts, for a total of 11.5 kWatts. Each sector has three closed loops that cool the preamps and interface boards. Water flow to the sectors is regulated and monitored through four distribution panels positioned on the detector carriage. The cooling water flow is interlocked, so that the power cannot be on unless the cooling water is flowing. The cooling system has been operated without interruption for 3 months during the commissioning period. Flow to the sectors has been balanced and has provided the necessary cooling and stability to operate the EMCal electronics.

The EMCal also uses four additional water chillers, located on the chiller platform in the Assembly Hall, to more precisely control the temperature of the SiPM optical sensors. These devices are not expected to generate much heat, but their gain and noise characteristics have a temperature dependence, and active cooling is expected to help compensate for the effects of radiation damage over time. Four insulated distribution panels, positioned on the carriage, are used to regulate and monitor the SiPM cooling water temperatures and flow rates. During commissioning, the chillers were operated at 16°C . The system was successfully balanced and operated to equalize the range of SiPM temperatures to within approximately $\pm 1^{\circ}\text{C}$ and maintain stable temperatures over time. A snapshot of the distribution of measured EMCal SiPM temperatures is shown in Fig. 35.

The eight cooling water distribution panels were designed, assembled, tested, and installed after the installation of the EMCal detector. These panels also have numerous safety/flow interlocks that needed to be tested and commissioned before use. A nitrogen distribution system also delivers dry gas to each of the sector enclosures to minimize moisture and condensation. The moisture level inside the sectors is monitored by a system of temperature/humidity probes, which are read out and monitored by the shift crew. The nitrogen flow has been balanced and moisture levels in the sectors have been maintained at minimum levels.

8.1.3 Post-Installation Commissioning

During the installation process, a mobile test stand was connected to each of the sectors in order to confirm functionality. Through this test stand, LEDs and electronic test-pulses were used to check the functionality of each channel in every sector at a stage where repairs would be doable.

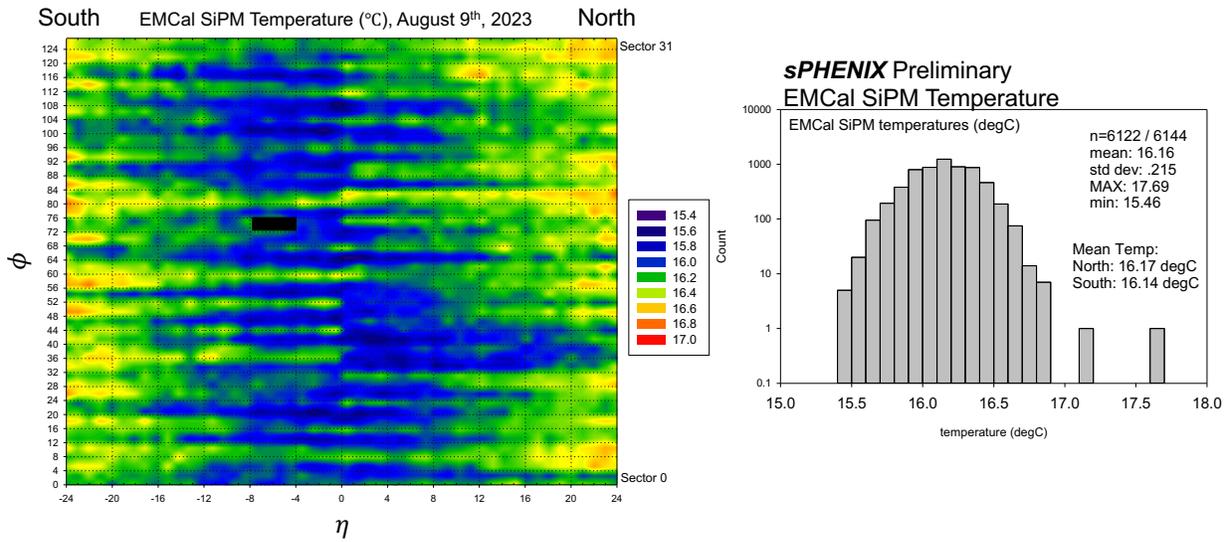


Figure 35: Distribution of measured EMCal SiPM temperatures throughout the EMCal detector.

At the end of the installation process, it was confirmed that only one of the 24,576 channels was behaving sub-optimally.

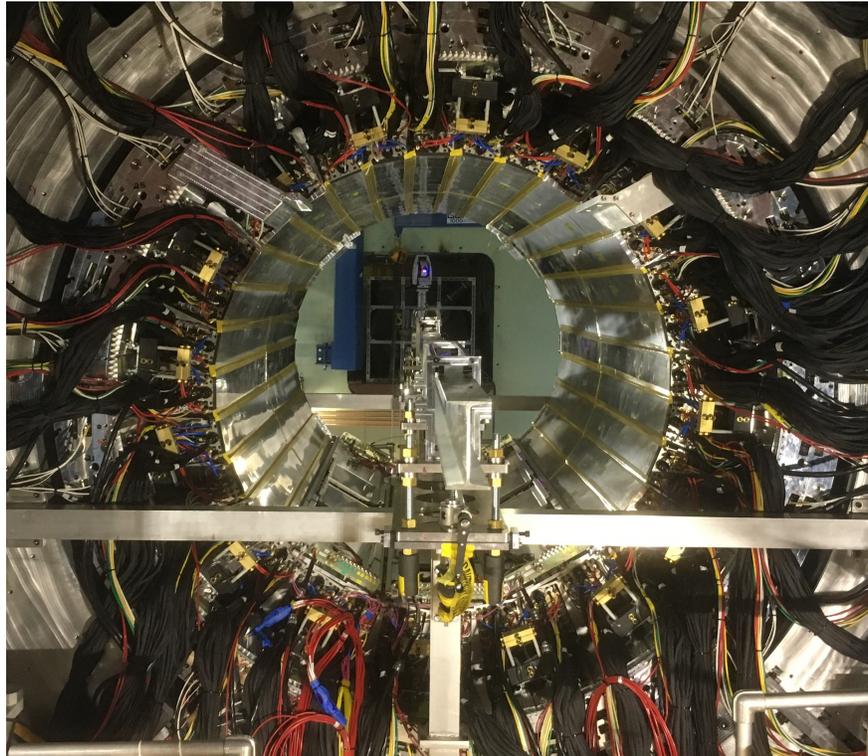


Figure 36: EMCal sectors (and TPOT) installed, with cabling and cooling manifolds, prior to installation of the TPC. The structure in the center of the detector and its support are parts of the TPC installation fixture.

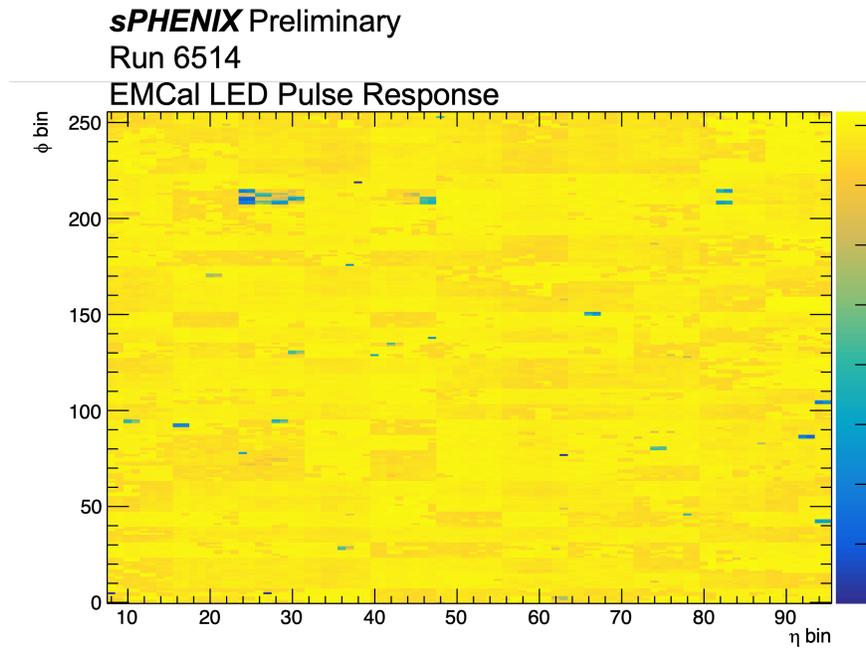


Figure 37: Plot showing the LED pulse response in the EMCal. Successful readout of >99.9% of channels where the EMCal is instrumented is observed. SiPM's in each channel are driven by an LED, whose pulse width is wide enough to drive it into saturation.

Fig. 36 shows all sectors fully installed and cabled and with the cooling manifolds installed. Once connected up to the full readout electronics in the racks, a series of pre-beam commissioning tests was performed in order to validate the functionality of the complete readout chain. Using a set of test pulses, LED runs, and pedestal studies, the functionality of the completely installed EMCal and readout was thoroughly tested in its final assembled state. The results of these tests showed that it was possible to read out the entire calorimeter using the full chain, and over 99.9% of channels were operating as intended – as shown in Fig. 37.

A considerable effort was also put into developing and testing control and monitoring GUI's for detector experts and the shift crews, and the procedures and instructions for using them. These GUI's are used to monitor the status of the EMCal, turn power ON/OFF, reset trips, monitor currents, temperatures, and humidity. Feedback from the shift crews was helpful in developing features, functionality, and documentation and making the GUIs useful and understandable to non-expert users.

Using beam data, the detector performance was demonstrated by the reconstruction of the diphoton mass distribution with the observation of a first π^0 mass peak. The resulting distribution for a short early run of 56,000 events can be seen in Fig. 38, where a peak is observed at approximately 90 ADC. Through the combination of the various commissioning stages the EMCal has been successfully commissioned as a functioning electromagnetic calorimeter with over 99% live channels.

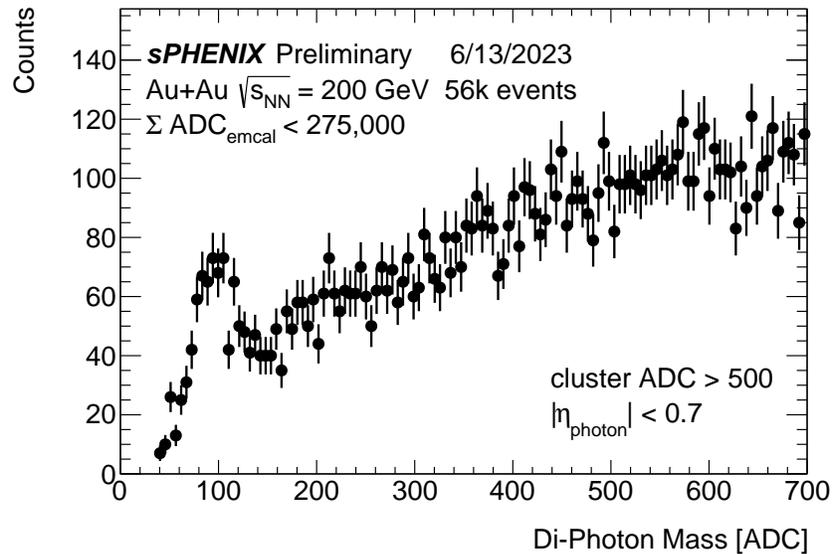


Figure 38: Shown is the di-photon mass distribution in units of ADC showing a clear π^0 mass peak

8.1.4 Leakage Currents

In an effort to monitor the health of the electronics, monitoring was developed to record the leakage currents in the EMCal. Fig. 39 shows the leakage current as a function of the number of weeks into the run (i.e., as a function of time). One can see it is steady for the first two weeks with no beam, and then begins to jump and increase steadily after beam is first turned on. The gap in coverage around week 11 is due to the fact that the database that handles these measurements crashed, and thus the data was lost. The leakage current was less than 2% of the maximum supply current at the end of the run.

8.1.5 Remaining Commissioning Tasks

Additional statistics, beyond those available in the 2023 data, are required to demonstrate the overall energy resolution and uniformity of the EMCal response. With the limited commissioning time that was available due to the shortened 2023 run, it was not possible to commission the EMCal for the full 15 kHz running. With the current readout setup, we were able to achieve data taking rate of approximately 2 kHz. Through the implementation of a new data compressed format, as well as the inclusion of zero suppression, the experiment is well on track to be able to record EMCal data at the 15 kHz target.

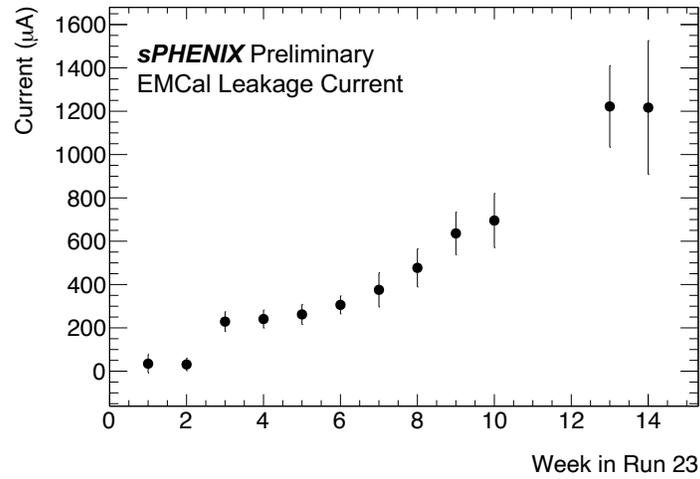


Figure 39: Weekly leakage current recorded over the course of Run 23. The current is the average of 64 towers (256 SiPM's) served by one Interface Board, and the error bar is the standard deviation over the 64 towers. One can see an initial bump at week 3 when the beam is turned on. The gap around week 11 is the result of a database loss.

8.2 Hadronic Calorimeter

The hadronic calorimeter (HCal) in sPHENIX has two layers, the inner HCal and outer HCal, just inside the magnet (116-137 cm radially) and outside the magnet (182-269 cm radially), respectively. Both HCals have a tilted plate design with scintillator tiles sandwiched between the plates. The tilt angles are determined such that a straight track would intersect 4 tiles in both layers. The scintillator tiles have an embedded wavelength shifting fiber which sends light to a silicon photomultiplier (SiPM). Signals from the SiPMs are grouped together with a preamp to define a tower which is then read out through the rest of the electronics. The HCal plates were assembled as 32 sectors in azimuth (ϕ) that span the full η range. Each sector contains 2 towers in ϕ and 24 towers in η . Therefore each HCal contains 64 towers in ϕ and 24 in η . There is a one to one correspondence between tower and read-out channel.

There are a few differences between the two HCals. The plates are steel in the outer HCal and aluminum in the inner HCal. While the larger, denser outer HCal measures most of the energy of the hadrons, the inner HCal is important for measuring hadrons that start showering as they traverse the EMCal. Because the solenoid magnet is designed with a chimney needed for the cryo-system that sits within the acceptance of the outer HCal, the steel and tiles in this region are shorter. This region is 6 towers wide in ϕ and 4 towers wide in η .

Prior to construction, there were several prototype studies of the HCals at the Fermilab Test Beam Facility, which utilized steel plates in both HCals and explored the performance of the tiles at mid rapidity ($\eta=0$) and higher rapidity ($\eta \approx 1$) [6, 7]. The energy resolution measured in all configurations was found to be significantly better than the design specification of $100\%/\sqrt{E}$ required to achieve desired jet resolution. In addition, these studies demonstrated the excellent agreement between the beam test data and results from GEANT4 simulations.

During the assembly process, the performance of the HCal tiles was studied using cosmic rays, and the results serve as an initial calibration. A cosmic ray trigger implemented during the commissioning period - see Section 8.2.3 for details - enabled additional cosmic ray measurements with the final electronics over the life of the experiment. In addition, all tiles can be illuminated and monitored with LEDs. The following text describes the commissioning achievements of the HCals regarding the assessment of live channels, implementation of readout with the data acquisition system, and covers data collected with collisions, cosmic rays, and LEDs.

8.2.1 Pre-Beam Commissioning

Every assembled HCal sector was tested three times: after assembly; pre-installation in the experiment; and post-installation. This ensured that all HCal channels were functional at every stage. Any faulty channels found were repaired by swapping components. Three tests were performed: a test charge injected into the pre-amplifier (per channel); an LED test (per SiPM); and a cable routing test. The test pulse is the simplest test as it does not require bias voltage being applied to the SiPMs. It tests the read-out chain all the way to the pre-amplifier boards, but not to the SiPMs. The LED test injects light into individual tiles and tests the entire chain all the way to the SiPMs. Employing these two separate tests allows to identify the origin of any malfunctioning components. The cable routing test works by reducing the bias voltage of individual channels sufficiently to effectively switch them off. This is done through a bias offset

that can be applied to individual channels through the cable that combines low voltage, bias, test pulse, and communication (power cable). It is then checked if the correct channel appears in the read-out. If either the power or signal cable is accidentally mis-routed, the wrong channel will appear. And mis-routings were identified and corrected.

In addition to these tests, one hour of cosmic data was collected for each sector after assembly (12 hours for the inner HCal due to the smaller tower size). The data collection was facilitated by the self-trigger capability of the digitizers used for the post-assembly tests and calibration data taking, which stem from the PHENIX HBD detector (i.e., not the sPHENIX production digitizers). Additional cosmic ray studies were done with sectors rotated in various positions as well as sandwiching individual towers between external triggers.

8.2.2 Data Acquisition, Sampling and Triggers

Data from the HCal is read out to Analog Digital Converters (ADCs), providing digitized data in 12, 16 or 31 programmable time samples for each tower in each event. This data is then transferred on GL1-accepted events to DCM-II and then PAR-III modules, and finally processed by jSEB-II cards, and stored as data files as is detailed in Section 6.

The HCal has been uniquely able to test the capacities of this system, as a smaller computational scale than that of the EMCal and as the earliest installed system. To this end, the HCal commissioning has been tied to that of the DAQ system, marking milestones of being able to read out a single crate of digitizers (corresponding to a quarter of all channels), to then a full rack (corresponding to a single DCM), and finally the whole detector. Until the middle of July, the HCal was taking 31 time samples for each event; however, all the necessary components are in place to allow us to switch to 12 or 16 time samples per event, which recent tests with beam triggers and clock triggers showed gave a factor of 2 increase in speed, to $\sim 3.3 - 3.6$ kHz rates at 16 time samples and ~ 4.2 kHz with 12 time samples. This is below the target rate of 15 kHz, however, this is not using multi-event buffering, zero suppression or modifying dead time, all of which are expected to contribute further factors of 2, which would put us well within our target rate. Implementation of these changes is discussed in Section 6, as this is a development in the DAQ that impacts a wide range of subsystems. The HCal will continue to use clock triggered, random triggered and pulsed LEDs to test timing and data acquisition efficiency as upgrades are made in order to be able to compare known rates across versions.

8.2.3 HCal Trigger

The HCal will provide trigger primitives for the jet trigger to be used during proton-proton collisions, as described in Section 7. These consist of the energy sum of 2×2 towers in the HCal, corresponding to an area of 0.2×0.2 in $\eta - \phi$ space. These trigger primitives are currently implemented for triggering on cosmic rays in the outer HCal. The trigger requires a 2×2 tower sum above 640 ADC within the outer HCal. Between stores in RHIC, data is routinely collected using this cosmic ray trigger for use in calibrating the HCal.

8.2.4 Cosmics

The main objective of the cosmic ray studies with the sPHENIX operating setup for the HCals, is to perform a comprehensive tower-by-tower absolute calibration to the electromagnetic energy scale, by comparing the data to GEANT4 simulations. Relative calibrations derived from the HCal assembly test stand data in building 912 is applied as a preliminary calibration to reduce the possible bias from the absence of the calibration. A quality assurance test using only the relative calibrations and overall energy scale is used to validate the simulation.

8.2.5 Tower by Tower Absolute Calibration Derivation

The trigger used for obtaining cosmic data is described in Section 8.2.3, the trigger threshold was set to 0xA which corresponds to 640 ADC count for the 2 by 2 tower sum. Notably, typical cosmic muons generate a signal approximating 2000 ADC in the oHCal, with pedestal fluctuations around 10 ADC. Consequently, the threshold settings can eliminate the impact of pedestal fluctuations and capture the majority of cosmic muon events.

The raw cosmic ADC signals initially undergo first-order relative calibration using the factor `rel_calib`, derived from the cosmic data taken in building 912, which results from the cosmic muons traversing all tiles in the tower. An additional correction is applied to eliminate the temperature influence in building 912:

$$\text{rel_calib} = N * \text{MPV}_{\text{simulated}} / \text{MPV}_{912} \quad (1)$$

where N is the normalization factor to bring the mean to 1, implying that this calibration will not induce a shift in the mean MPV position for the sPHENIX cosmic data.

This factor primarily eliminates the performance differences between towers, except for the differences attributed to tile shapes (η dependence) and temperature variations in 1008 (time dependence). The calibration function is expressed as

$$\text{tower_ADC} = \text{tower_ADC}_{\text{raw}} * \text{rel_calib} \quad (2)$$

An offline cut, designed to select cosmic muons that traverse all tiles within a tower, maintaining an approximately perpendicular trajectory to the detector barrel's radius is applied. This step is critical for quantifying the overall performance of each tower, which encompasses all the tiles in the given tower.

With these initial calibrations and cuts in place, we are able to determine the ADC distribution of cosmic muons for each tower. A comparative figure of the sPHENIX data and the corresponding GEANT4 simulation data is depicted in Fig. 40. The observed agreement validates the reliability of these two data sets and there are no obvious biases in the data processing and analysis methods. A Gamma function fit is then applied to this distribution, thereby obtaining the Most Probable Value (MPV) for each tower. Once this fitting has been applied to both the sPHENIX cosmic data and the simulated cosmic data, we can represent the general performance distribution for all towers in the 2D distribution, as depicted in Fig. 41. The similarity in the MPV structure seen in the 2D distributions is observed between the sPHENIX experimental data and the GEANT4 simulation data, which illustrates the reliability of the result, and this allows for the determination

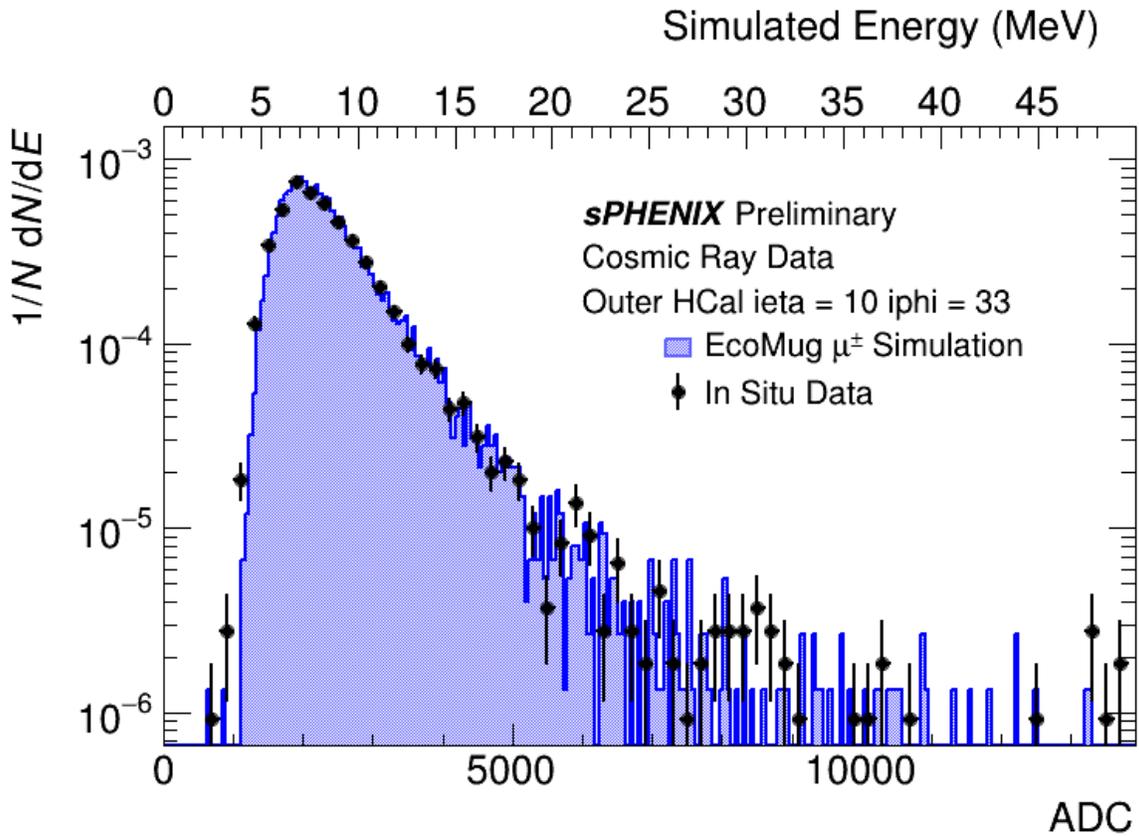


Figure 40: A comparative figure of the sPHENIX tower's energy distribution (solid points) and the corresponding GEANT4 simulation result (blue shaded region).

of an energy scale from these two datasets. The observed structures in these distributions are attributed primarily to geometric effects due to changing the orientation of the incident cosmic rays. The region in the lower ieta range corresponding to $i\phi = 16 \sim 18$ represents the chimney towers for oHCal, characterized by shorter tiles. These towers have a lower statistics and pose greater fitting challenges. When taking the ratio of these plots, the variation for a majority of the detector is less than 2%. However, a clear temperature dependent effect has been observed with higher temperatures observed for the upper portion of the detector. A correction for this temperature effect has not yet been applied to the data and the deviations observed are within the expected size of the correction.

This outcome agrees with the expectations of the application of the preliminary building 912 calibration factor, designed to eliminate all influences except those attributable to temperature. An ongoing temperature correction study aims to restore the gain shift observed in the ϕ direction, obtained from the sPHENIX LED data.

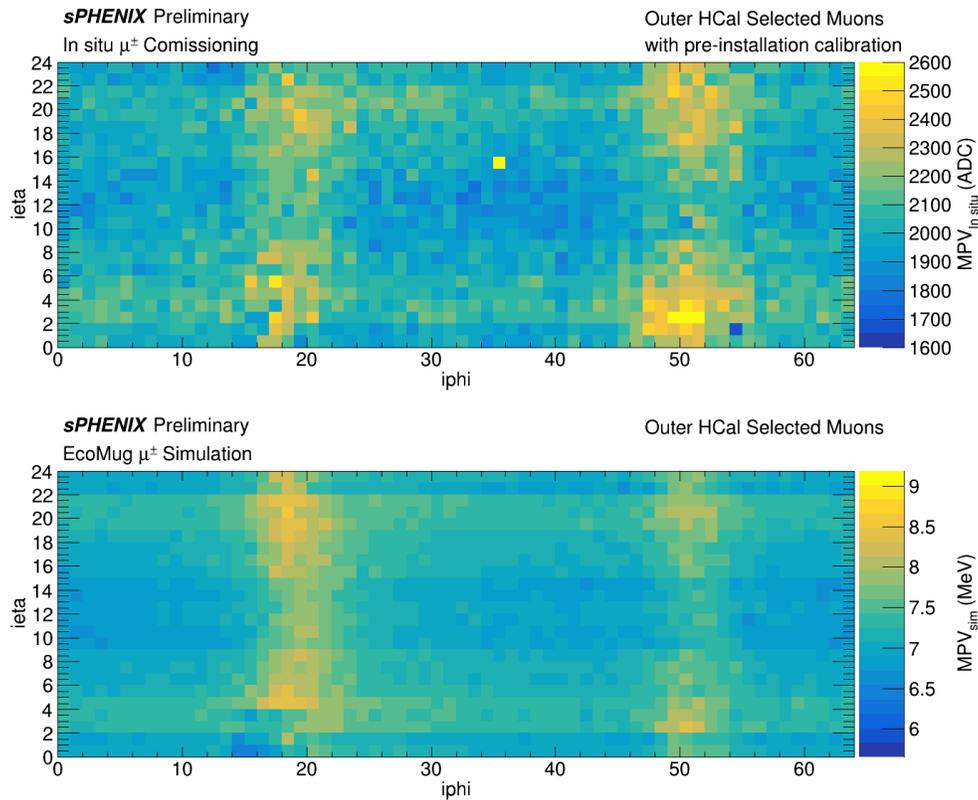


Figure 41: The 2D MPV distribution for the sPHENIX cosmic data (left) and the simulation data (right).

8.2.6 LED

To monitor the impact of radiation damage and to search for hot towers in the HCal, we take an LED data run of 10k events once every two hours while running physics data or cosmics data. These runs give us a known response with a distinguishable, uniform waveform signal in every tower of the HCal as is seen in Fig. 42, and, by monitoring the peak height, peak width and pedestal RMS over time, give a proxy for radiation damage on an SiPM by SiPM basis, as well as a known response for testing HCal response to system changes, calibration for temperature and gain changes, and identification of dead channels or hot channels that develop during the process of the run. From monitoring the health of the detector in both LED and collision data, we conclude that overall all channels are working well. Occasionally a hot tower appears, but can usually be addressed by re-initializing the ADCs, i.e. those hot towers are not related to the detector itself and can be easily fixed. In rare cases, hot towers have appeared that were related to a faulty pre-amplifier. We were able to fix these in all cases by changing the pre-amplifier board.

In addition, we continue to monitor the leakage current and have found that there was an increase in the rate of increase of leakage current as the beam increased in intensity, as expected. However, the change has remained linear and still has a projected total increase in leakage current within acceptable limits. This is reflected in Fig. 43

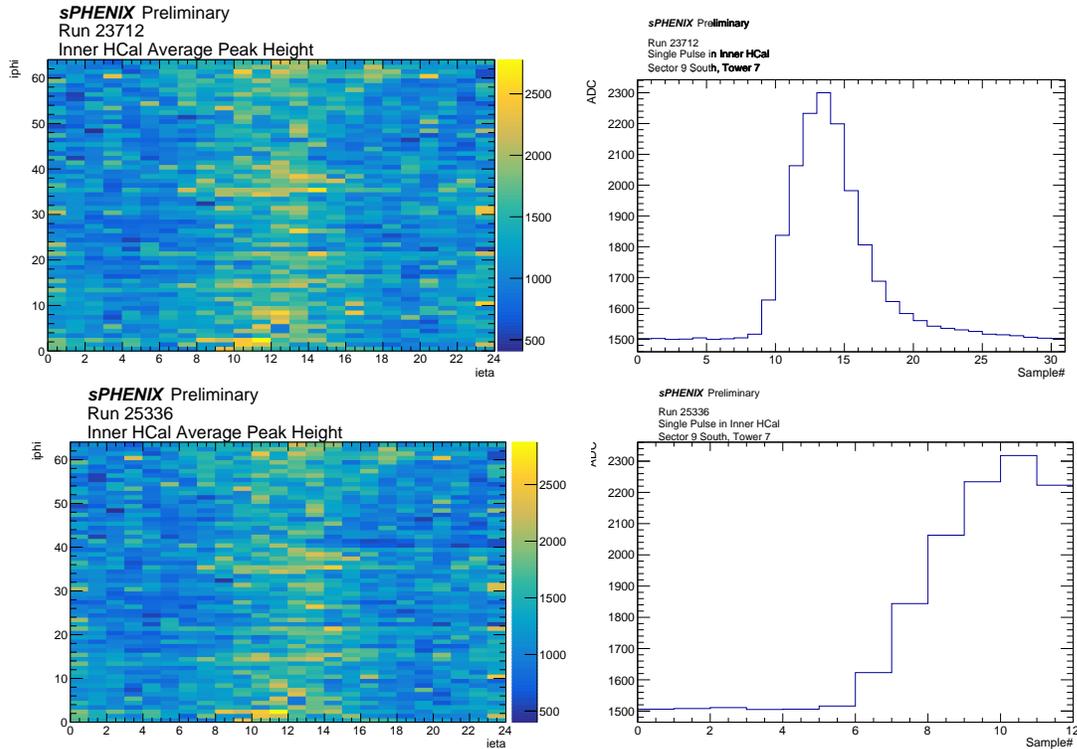


Figure 42: Top: Inner HCal LED response with 31 samples (Left) Average Peak Height in all sectors across 10,000 events. (Right) Single event waveform showing ADC response in each time sample. Bottom: Inner HCal LED response with 12 samples. (Left) Average Peak Height in all sectors across 10,000 events. (Right) Single event waveform showing ADC response in each time sample in same channel as top right.

8.2.7 Correlations with other subsystems

Due to the independent readout of each detector subsystem, the full event building of sPHENIX is done offline using information from the various front-end electronics systems. The HCals are read out in two output streams, corresponding to the East and West halves of the system. These two streams correspond to two Sub-Event Buffers (SEBs), each of which provides eight “packets” of data. One packet corresponds to the readout of four HCal sectors. For each packet, an event number and beam clock time are recorded for use in event synchronization. This offline event building can be validated by analysis of the correlations between separate subsystems.

Figure 44 shows the correlation between the total energy in the inner and outer HCals. The energy in each tower is determined using a simple waveform processing which subtracts the pedestal, determined using the average of the ADC readings in the first three time samples, from the peak of the waveform. Energy deposits from noise and out of time backgrounds are suppressed using a timing cut, where only towers with an in-time peak are selected. The axes are arbitrarily scaled to range from 0 to 1. The expected strong correlation between the inner and outer HCals, which cover the same $\eta - \phi$ space, is shown. This correlation confirms the proper alignment of the 16 packets of information read out from the HCals. Figure 44 also shows the correlation between the total energy in the outer HCal and the total charge in the MBD. These two detector systems

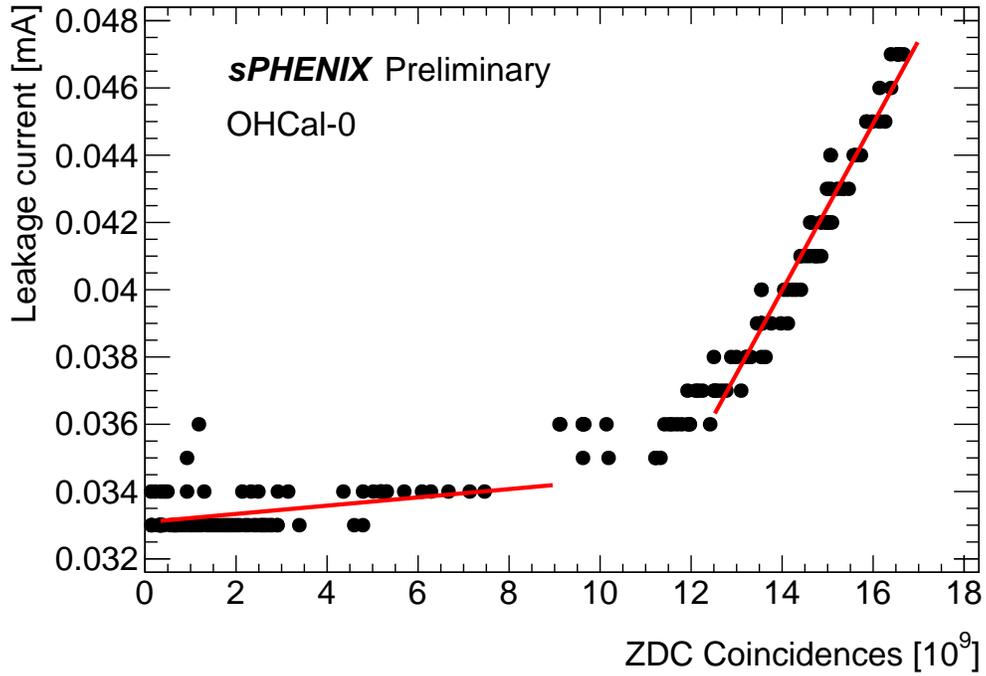


Figure 43: Leakage current in HCal measured once per fill as a function of total number of ZDC coincidence hits

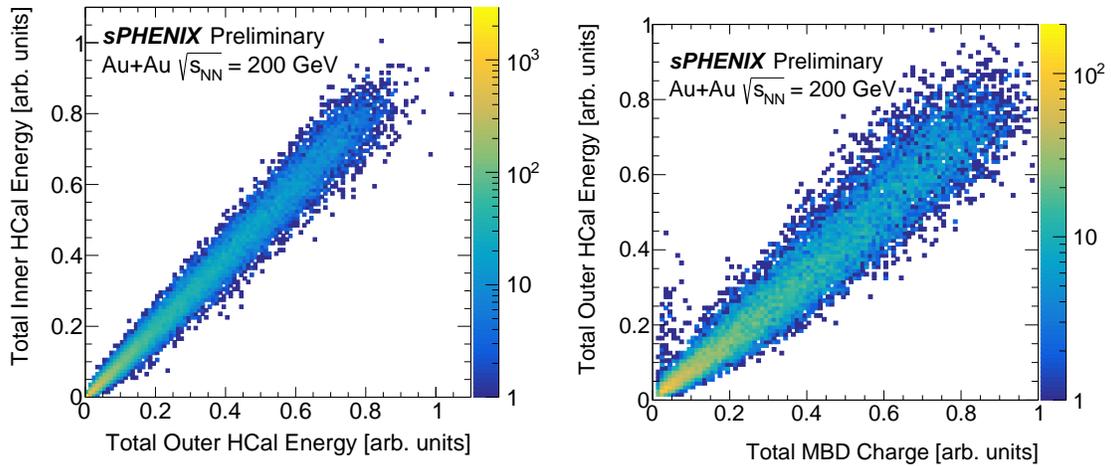


Figure 44: The correlation between the total amount of energy in the inner and outer hadronic calorimeters (left) and the total amount of energy in the outer hadronic calorimeter with the total amount of charge in the MBD (right) in Au+Au collisions at $\sqrt{s_{NN}} = 200$ GeV.

cover distinct areas in η . The correlation again shows the ability to align the two detector systems offline using information from their front-end electronics.

8.2.8 Remaining commissioning tasks

The previous sections highlighted the successful commissioning progress achieved during the 2023 RHIC run. However, commissioning tasks that could not be accomplished before the unexpected end of the 2023 RHIC run remain.

- Analyze cosmic ray data for each tower in both the inner and outer HCals. This requires sufficient cosmic ray data for all sectors which requires a large amount time for some sectors due to their orientations with respect to the incoming cosmic rays. A large sample of cosmic ray data is being collected following the conclusion of beam running.
- Apply temperature corrections based on LED data to the calibrations extracted from the cosmic ray studies. Initial studies related to this are currently underway but the process of how this correction will be applied during the run will not be fully exercised until the beam returns.
- Regularly exercise the high rate DAQ capabilities with the HCals. Since the commissioning period of the 2023 RHIC run was cut short, we did not completely shift to a regular data taking mode. Therefore the 15 kHz data taking rate goal will need to be achieved in the next RHIC run.

9 Tracking Detectors

9.1 MVTX

The monolithic active pixel vertex detector (MVTX) is the innermost detector at sPHENIX and is designed for precise primary and secondary vertex determination. It covers a radial distance between 2.4 and 4.8 cm from the nominal IP and has a z coverage of ± 13.5 cm. It consists of three layers of sensors and uses the ALPIDE chip for particle detection [8] which has a hit resolution of approximately $5\ \mu\text{m}$ and a tunable integration time which is typically between 5 and $10\ \mu\text{s}$. Each ALPIDE chip is $1.5\ \text{cm} \times 3.0\ \text{cm}$ in size with a pixel pitch of $29\ \mu\text{m} \times 27\ \mu\text{m}$ and a pixel matrix that is 512×1024 pixels. The in-pixel circuitry consists of amplification, pulse shaping, discrimination, digitisation and a multi-event buffer (MEB) with a depth of 3 events. The ALPIDE chips are arranged on a “stave” which is the basic building block of the MVTX. Each stave consists of 9 ALPIDE chips, a digital power line, an analogue power line, a clock line, a control line, nine high speed signal lines, water cooling lines and a carbon fibre wire frame for mechanical support. The chips, power and signal lines are bonded to a flexible printed circuit (FPC). As the material budget of a vertex detector is always a concern in high energy physics experiments, on average, a stave is 0.35% of a radiation length ($0.0035 \times X_0$). The inner layer of the MVTX (L0) contains 12 staves, the middle layer (L1) contains 16 staves and the outer layer (L2) contains 20 staves. This results in a total of 226,492,416 pixels. The MVTX was built using so-called “gold” staves, meaning that each stave has a low number of dead, noisy or stuck pixels which means that over 99.99% of pixels are capable of being read out.

The staves are read out at a rate of 600 Mbit/s or 1.2 Gbit/s via 28 AWG high speed cables to the off-detector electronics. Each staves' off-detector electronics comprises of a Readout Unit (RU) and a Power Board (PB). The ALPIDEs send data to the RU's using 8-bit/10-bit encoding. The RU packages the data from all 9 chips and sends the data to a FELIX card over 3 optical links using the GBT protocol [9]. The PB provides analogue and digital power to the stave circuitry. A PLC based safety system monitors the stave temperature from a PT100 attached to the FPC. The data from eight staves is combined at a FELIX card, and the individual data from each RU are then packaged and sent to the data acquisition system running on the host servers over 16-lane PCIe connections.

9.1.1 Installation

The MVTX supporting structure was installed on the 27th of March 2023 and surveyed into position on the 28th of March. This supporting structure is connected to the frame of the inner hadronic calorimeter at four points and helps define the final position of the MVTX. From earlier surveys it was determined that the inner HCal was installed 7 mm south of the nominal IP and the INTT was installed 5 mm south of the nominal IP. As the z-separation of the closest points of the MVTX and INTT is 5 mm, which was accounted for in the design of the supporting structure, this results in 2 mm of extra clearance between the MVTX and INTT and thus no extra adjustments of the supporting structure were required.

As the MVTX is cantilevered from the supporting structure on the South side, it was important to ensure that the mounting point of the MVTX to this structure is as co-planar with the vertical axis as possible. The separation between the point that the MVTX mounts to the supporting structure¹ and the end of the detector is approximately 2m while the separation of the beam pipe and the MVTX was designed to be 2 mm with a 0.5 mm tolerance. This results in a maximum deviation of 0.25 mrad of the support from the vertical axis.

The MVTX was successfully installed on the March 30, 2023, symmetrically around the beam pipe. After installation, each stave was tested and found to still be classified as "gold." Due to the discovery and replacement of magnetic screws for the cable mounting points the MVTX was fully cabled on April 3, 2023 in preparation for commissioning. In the days following the installation, it was noted that the MVTX had come in to contact with the beam pipe. Sagging was anticipated from the metrological studies of the detector but it was uncertain that this contact was caused by the sag or if other actions around the detector had caused the beam pipe to rise. The closest beam pipe support point to the IP comes from the "spider" at the north end. This spider consists of four threaded rods, two in the vertical and two in the horizontal axis, around a peak clamp. The beam pipe was lowered using the vertical rods by 2 mm and the MVTX was not observed to travel with the pipe. The MVTX was monitored daily for the next 10 days and no further sag was noticed. Pictures of the beam pipe before and after the change in the beam pipe position can be seen in Fig. 45.

¹The mounting point of the MVTX is a separate ring and hence can be manipulated independently from the position of the general supporting structure

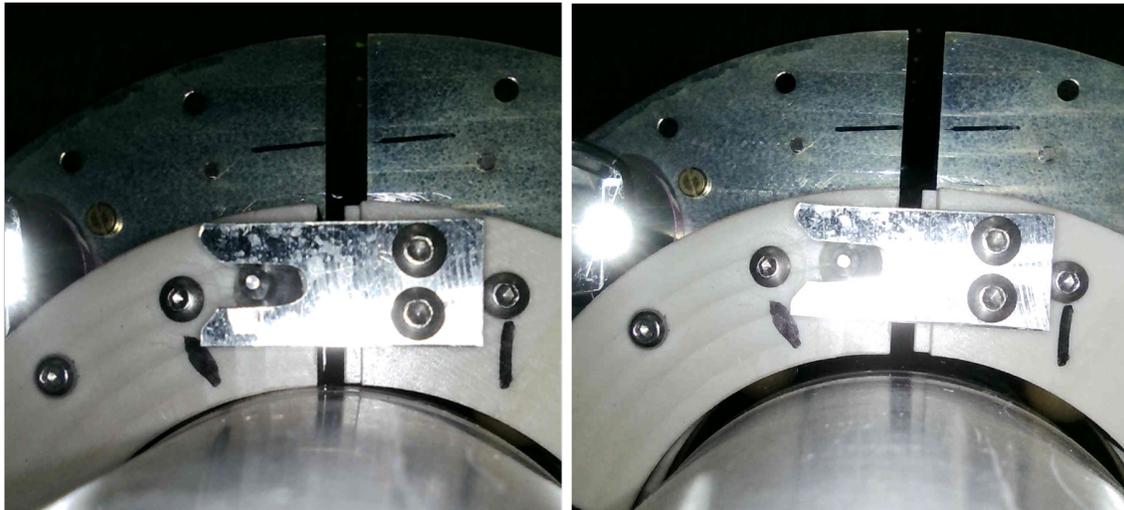


Figure 45: Photographs of the MVTX after installation taken from the north end of sPHENIX. The MVTX is the structure with the vertical split while the beam pipe is the lower object. These pictures were taken before (left) and after (right) the beam pipe was moved 2 mm lower in the vertical axis.

9.1.2 Cooling

The MVTX is cooled using distilled water treated with algacide (glutaraldehyde), and algae growth and corrosion inhibitor (Optishield+²). Both the MVTX detector and the off-detector readout electronics (the readout units and power units) are cooled using subatmospheric chillers provided by Chilldyne called the CDU15 (Cooling Distribution Unit)³. The staves require one CDU15 with a flow rate above 1.4 L/min while the readout electronics require two CDU15s, each with a flow rate above 6.0 L/min. The readout electronics are housed in 4 crates in a single rack and one CDU15 supplies cooling to the top two racks while the other supplies the bottom two racks. Custom built distribution panels are used for each CDU15 and are located on the carriage close to the racks and each panel sends information on the input and output temperatures and pressures which are monitored in the control room. The stave cooling panel also contains flow switches which are also monitored in the control room. Each CDU15 sends information on its flow, internal temperatures and status codes which are monitored by the shift crew. The MVTX chiller system has no interlock, however each individual stave and readout unit monitors its own temperatures and will interlock sectors of the MVTX if any temperature goes above a predefined limit (30°C for a stave and 45°C for a readout or power unit). The CDU15 error and flow rates are programmed to issue audible alarms in the control room if they go outside normal operating conditions, and all CDU flows and status are recorded by shift crews 6 times a day in addition to continual data logging.

During detector operations in May, it was noticed that the flow rates of the readout electronics chillers was decreasing with time. This could not be remedied by evacuating air in the system and the rate of decrease was slow (1 L/min or so over a week) so it was not believed to be due to air build up or leaks in the system. On June 12, 2023, the filters on the system were inspected and

²<https://shop.optitemp.com/>

³<https://chillydyne.com/>

found to have significant build up of residue so each filter was replaced on these two systems, totalling 24 filters. After 3 weeks, the filters were inspected and no further residue was noticed. The flow rate also increased above previous levels and has been running at a constant flow rate since the replacement. The last replacement of these filters was performed in August 2022 so it has been decided to inspect and replace these filters before the start of each run. A photograph of the filters before replacement can be seen in Fig. 46. The origin of the residue has not yet been determined as to whether it came from dirt and dust in the cooling tubes, hoses and pipes or whether it was bacterial growth. It should be noted that, in 2022, there were periods where no algaecide was used in the system and less algae growth inhibitor was added to the cooling water based on earlier treatment procedures provided by the manufacturer.

Since replacement of the filters, the MVTX cooling system has operated with a flow rate of approximately 1.6 L/min for the staves, and 8.4 and 9.0 L/min for the two readout electronic CDU₁₅'s.

In addition to the cooling system, the MVTX staves are kept in a nitrogen-rich environment using 12 tubes with an ID of 3 mm to avoid corrosion. The flow rate of nitrogen to the MVTX is kept at approximately 1.5 scfm.



Figure 46: Photograph of an MVTX readout electronics cooling filter from 12th of June before replacement.

9.1.3 Slow controls and monitoring

The MVTX is controlled using WinCC software ⁴. The control software was initially developed by the ALICE collaboration at CERN to control their vertex detector, the ITS-2, of which the MVTX is derived. The control software has been adapted for use with the MVTX. A single computer is used to issue commands to all 6 MVTX FELIX servers in parallel which then distributes commands to the readout units, power units or staves as required. The software is designed to only allow control by a single user at one time to avoid command conflicts. However, the MVTX control can

⁴<https://www.siemens.com/global/en/products/.../simatic-wince-v7.html>

be compartmentalised to subsystems as needed so, for example, if there is an error in a single stave, the reset commands can be issued to that single area without affecting the rest of the system. The WinCC package also controls the CAEN low voltage power system which supplies the bias to the readout and power units.

The underlying code that communicates to the detector is known as ITSCommLayer and is capable of issuing commands to configure the readout units, staves and perform voltage drop corrections amongst other features. During the commissioning period, the software was updated to allow global threshold settings to be applied to the MVTX and to load dead, noisy or hot pixels masks. Another feature enabled during the commissioning was an auto-recovery procedure. WinCC constantly monitors the status of the readout units for an indication of errors in the ALPIDE's or RU's and will then issue a recovery prompt as needed. Currently this procedure takes on average 12 seconds per stave as it will bring the entire stave offline and reset the readout unit and stave. There are several "wait" and commands that are not needed to recover the staves and work is ongoing to reduce the recovery from 12 seconds to between one and three seconds.

Monitoring of the MVTX is achieved using three software packages. WinCC, as mentioned above, is the control software for the MVTX and provides information to the shift crew on the cooling panel pressures and temperatures, the stave temperatures, the stave currents (both digital and analogue) and the auto-recovery status for each stave. The stave temperatures, cooling system pressures and temperatures as well as flow status are monitored by the PLC based MVTX safety system. This safety system is capable of turning off the CAEN power supplies in case of (programmable) out of bounds temperature conditions of the staves. The status of the safety system measurements is monitored and displayed by the WinCC software. The chiller flow rates and data acquisition status and rates are visualized by Grafana software. The chiller flow rates, temperatures and error codes are monitored by Ignition software in the control room. Examples of the WinCC and Grafana monitoring can be seen in Fig. 47.

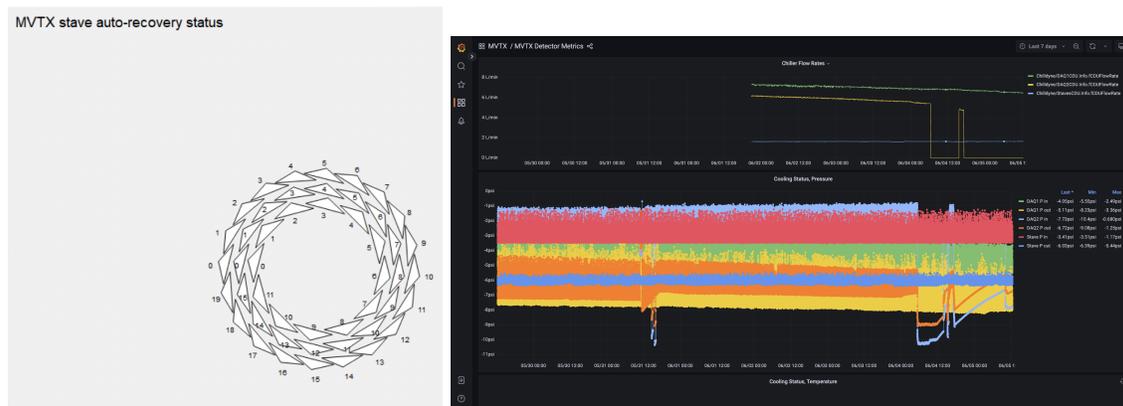


Figure 47: Monitoring of the MVTX in WinCC (left) and Grafana (right).

9.1.4 Calibration

The calibration of the MVTX is performed during no-beam periods in a two-step process. The first step is to perform a threshold tuning where the pixel configuration values can be altered sequentially for a whole ALPIDE. Each ALPIDE can have a different set of configuration values

but the values are the same for each pixel on that ALPIDE. Each pixel has a comparator circuit which checks if the charge collected is above a set value and a charge injector which can inject a programmable amount of charge to that pixel. For a desired threshold, two parameters, ITHR and VCASN, are altered until the pixel “fires” for that given charge injection. VCASN determines the threshold baseline and ITHR is the current that discharges the comparator. Thus, a higher VCASN or lower ITHR results in a lower number of electrons required to trigger a signal from the pixel. Four different electron values were tuned for; $100e^-$, $130e^-$, $150e^-$ and $250e^-$, with 100 electrons chosen as the default operating value.

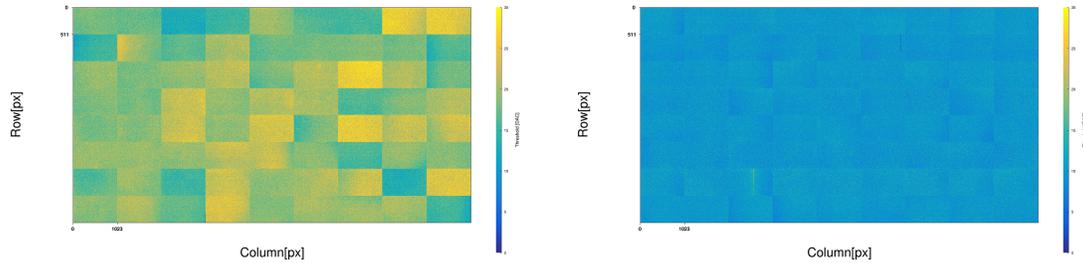


Figure 48: Pixel threshold distribution of all 8 staves connected to one MVTX Felix server, before (left) and after (right) threshold tuning calibration.

The second step in calibration of the MVTX is to take a “Fake Hit Rate Run” (FHR). An FHR applies the threshold values determined from the previous step to each ALPIDE and records which pixels fire more than a user defined value, typically more than once every 1 million triggers (readout frames). For a pixel to be read out, it needs to make a logical AND with a “strobe” pulse which is sent to all pixels simultaneously. The length of the strobe (and thus the length of the readout frame) can be set by the user, and during the commissioning period we used four values; 200 ns, 525 ns, 10 μ s and 89 μ s. Using a value for the strobe length larger than the ToT ensures that any pixels fired from cosmic rays can return below threshold before the next readout window. All recorded pixels are then written to a pixel mask and can be excluded from the readout during data taking. Using the threshold tuning and pixel masking, the MVTX noise rate can be brought to below 10^{-9} as seen in Fig. 49, much less than the design value of 10^{-6} .

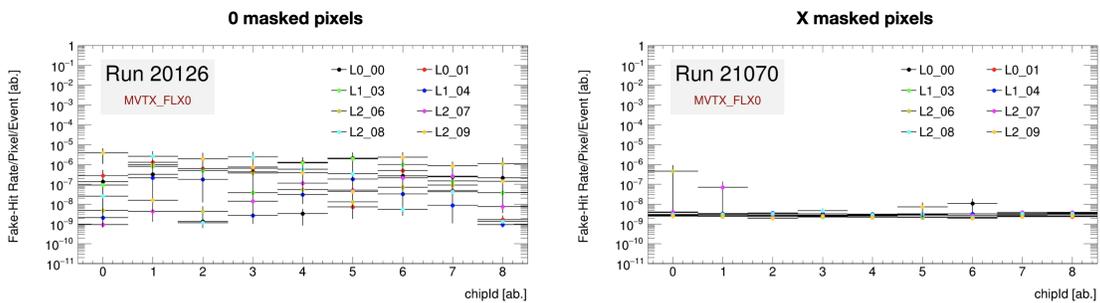


Figure 49: Noise rate recorded by one MVTX FELIX server before (left) and after (right) pixel masking as a function of chip ID.

After the thresholds and noisy pixel masks were determined, an initial sample of Au+Au data

was taken to study synchronisation and correlations within the MVTX. The detector is readout via 6 FELIX servers which each read 8 staves. The staves are distributed such that no FELIX reads out two staves that align azimuthally with the nominal IP to ensure that if one server stops reading data there are still two readout points for track reconstruction. The distribution of the MVTX staves by their software ID and associated FELIX card can be seen in Fig. 50. The internal correlation of staves on the MVTX from Au+Au collisions can be seen in Fig. 51. The data recorded by these staves comes from three different FELIX servers and so these plots show that the MVTX can be internally aligned. The broadening of the correlation between L0 02 and L1 03 is suspected to be due to beam backgrounds appearing in one staffe but not the other and hence are uncorrelated with the collision. The beam backgrounds are discussed in the next section.

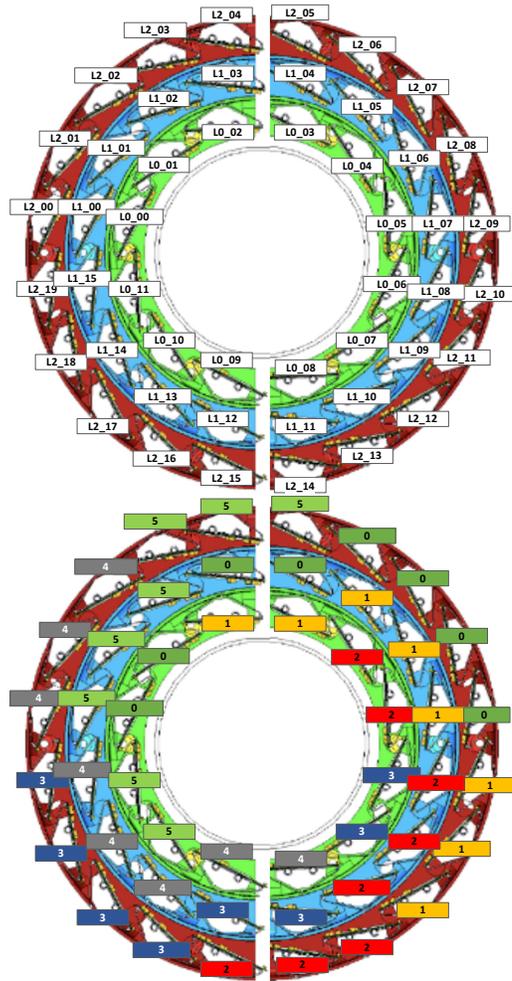


Figure 50: MVTX layout as viewed from the south of sPHENIX. Top - The software names of each staffe, bottom - the associated FELIX card each staffe is connected to.

9.1.5 Cosmic events

During periods without beam, it's possible to turn the MVTX on for data taking. As the MVTX is a streaming detector it will record continuously so since August 2023 we have run alongside the

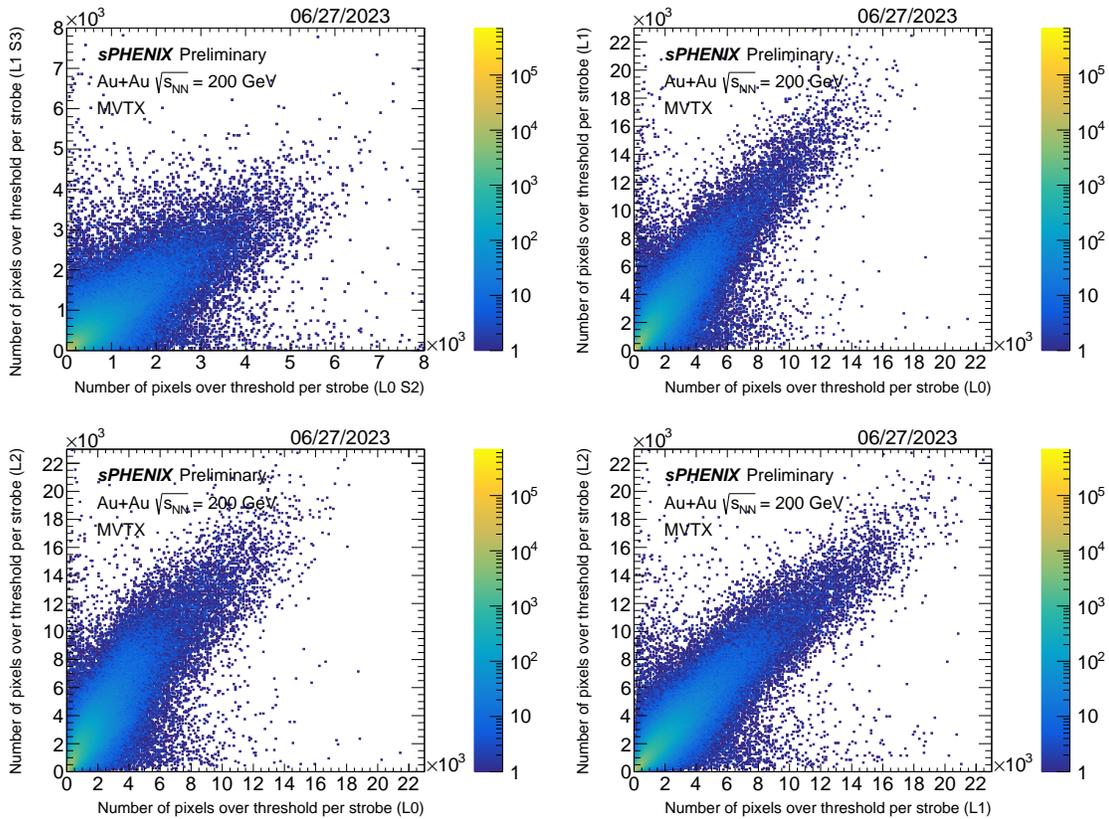


Figure 51: MVTX internal correlations. Top left - one stave on L0 vs one stave on L1 from different FELIX servers, top right - L0 vs L1, bottom left - L0 vs L2 and bottom right - L1 vs L2. These staves are distributed over three FELIX servers.

hadronic calorimeter which has a coincidence trigger to detect cosmic rays. The accept value of this coincidence trigger can be seen and tagged by the MVTX which is in turn used to filter out likely background events. The cosmic ray frequency for the MVTX is approximately 1 Hz while the streaming readout windows used for cosmic ray recording thus far are 500 ns and 89 μ s so this trigger is highly effective at reducing the search window for candidates.

The threshold was set to $100e^-$ and the obtained pixel mask from a fake hit run was applied to all sensors. The RHIC beam count orbit (BCO) was recorded in the data stream alongside the coincidence trigger and used for internal synchronisation of the FELIX servers and external synchronisation with other detectors. After decoding the data, the pixel hardware locations were translated to global coordinates in sPHENIX using the idealised GEANT4 geometry for the MVTX. From this stage, a simplified clustering algorithm was applied to ensure that a fired pixel had at least one adjacent fired pixel or else it was discarded. The validity of this simplified algorithm was confirmed by direct comparison with the clustering algorithm developed by the sPHENIX tracking group where the cluster positions agreed with each other. A cosmic ray candidate found in the MVTX can be seen in Fig. 52.

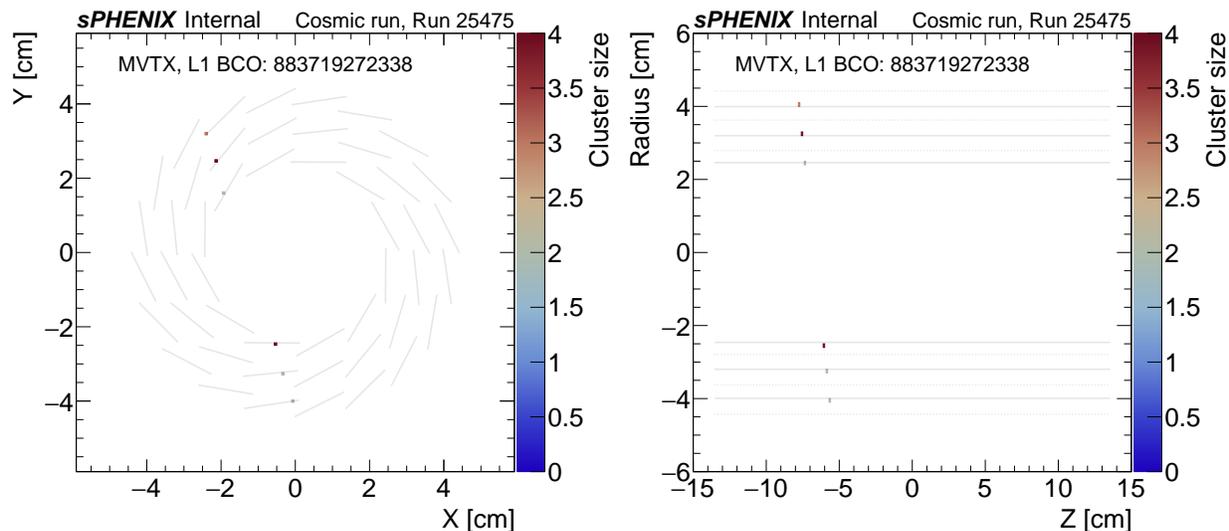


Figure 52: Two-dimensional event display of the event with BCO 883719272338: (Left) The cluster position and size in X-Y plane. (Right) The cluster position and size in Z-Radius plane.

9.1.6 Beam backgrounds

During initial operations, a larger beam background was observed in the MVTX. This background was large enough that it overwhelmed the readout which in turn disabled the readout of the affected ALPIDEs in the RUs after a few seconds of running. Inspecting hit maps of the MVTX from the initial period of a run revealed large "streaks" across the pixel matrix in the z-direction which is also the beam direction. These streaks were found to be composed of several hundred or thousands of pixels and the streaks were in the perpendicular direction to the readout and cross several chips which excludes the effect as being caused by the MVTX. The streaks appeared to have a radial dependence on their frequency (more in the inner layer) and not to occur across every stave for every event, leading to the conclusion of this being a beam background effect. A plot of the hit mask for a single stave in L1 on the horizontal axis for a recording time of 89 ms and 890 ms (one thousand and ten thousand readout strobes) with a single beam only in RHIC can be seen in Fig. 53. It can also be seen that between 89 ms and 890 ms the first chip of the stave stopped sending data due to errors.

Upon this observation, several tests were performed to either reduce the beam background from the accelerator or correct for the effects on the MVTX. From the MVTX, the threshold was increased from $100e^-$ up to $250e^-$, the strobe length was increased from $10\mu s$ to $89\mu s$, the detector was configured from triggered readout mode to continuous readout mode and an auto-recovery procedure was implemented in the slow control software. By changing from triggered to continuous readout mode, the behaviour of the pixel memory buffers is changed. In triggered readout mode, priority is given to the event that is currently being read out which means any incoming data is discarded after the buffer is full. In continuous mode, priority is given to the latest incoming data and so the event currently being read out is truncated when no new buffers are available. The pixel memory buffer is 3 events long. This change would mean that the pixel is cleared if the event takes too long to read out as would be the case in events were several

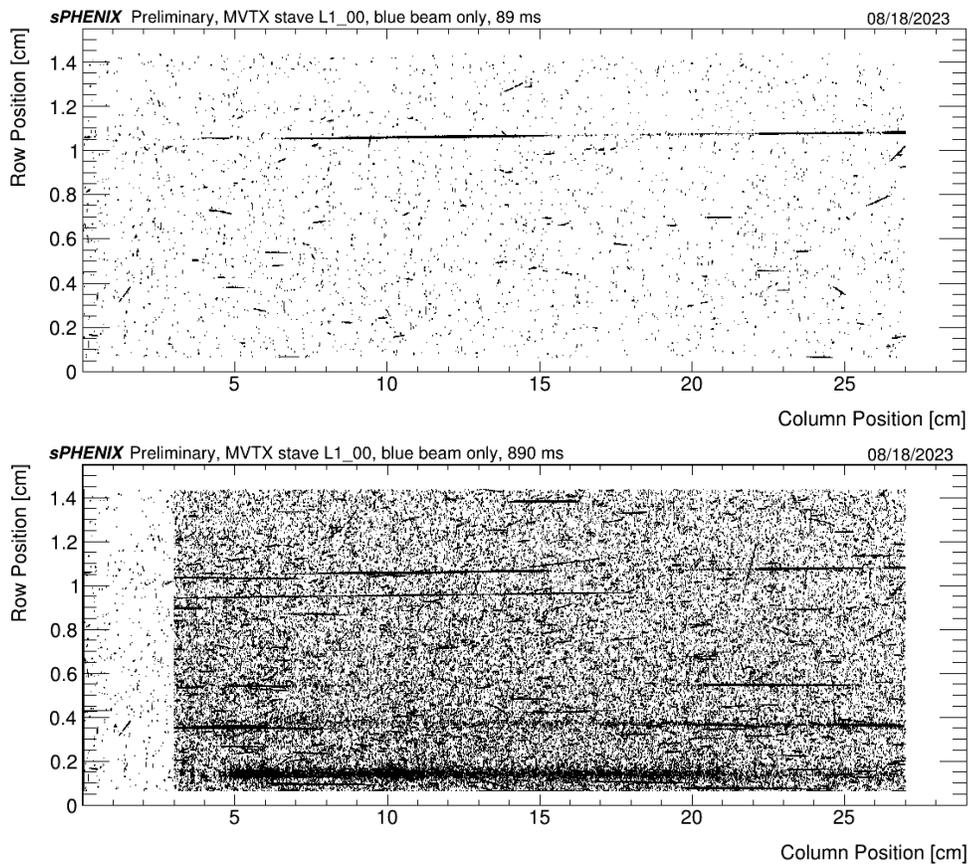


Figure 53: Hit map from Stave L1_00 taken during a fill with only the blue beam and no collisions. Top - data recorded over 89 ms. Bottom - data recorded over 890 ms.

thousands of pixels are fired on a chip. For comparison, it is expected that around 50 pixels would fire on a chip in a typical Au+Au event.

sPHENIX and the MVTX group worked closely with C-AD to understand the beam conditions and find parameters that would allow the MVTX to operate. The initial tests involved dumping one of the beams while holding the other in RHIC and observing the hit maps and error rates in the MVTX. It was found that in both the yellow and blue beams that the streaks and errors would still occur while the error rate seemed to be larger in the yellow beam compared to the blue. It was also observed that the error distribution had a preference to one side of the beam pipe but this preference changed for different fills. The MBD was also active for some of these tests and saw the same preference for increased hits in one side of their detector.

Another test involved decreasing the collimator gap for the yellow beam after the blue beam was dumped. The collimators are fully open during injection and close to their default value right before physics is declared. The values are based on beam studies to minimize beam loss while remove background effects and are dependant on the beam size. As the emittance of the beam decreases during a fill from the stochastic cooling and collisions, it is possible to close the collimators more. During this test, the collimators were moved after the beam emittance decreased from 3.0 mm-mrad to 0.56 mm-mrad which appears to be a plateau region reached a

few hours into the fill and can be seen in Fig. 54. The auto-recovery procedure was monitored and found to decrease from 179 recoveries, 388 calls to auto-recovery and 1589 instances of chips being in auto-recovery to 27 recoveries, 27 calls to auto-recovery and 61 instances of chips being in auto-recovery. It should be noted that the auto-recovery procedure takes 12 seconds and is queried every 15 seconds and so there are several cases where a chip will be recovering over two time periods and hence is double counted. A comparison of the staves that are in auto-recovery as a function of time, before and after moving the collimators can be seen in Fig. 55. Before moving the collimators in, the beam was moved vertically and horizontally to confirm it was in the optimal position. If was found that this was the case to within 1 mm in both the horizontal and vertical directions. This was confirmed by observing the beam background counters on either side of sPHENIX of which there are 6 in total and looking for the minimum background rate. It is possible that this test scraped more background events out and decreased the error rate in the test after moving the collimators.

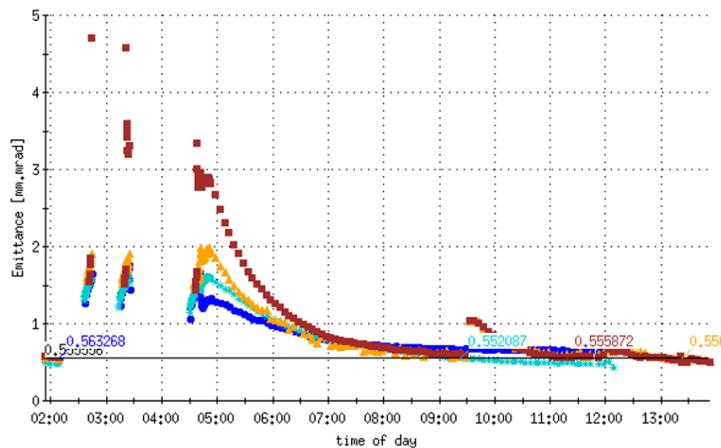


Figure 54: RHIC beam emittance during the fill in which the collimators were moved closer together.

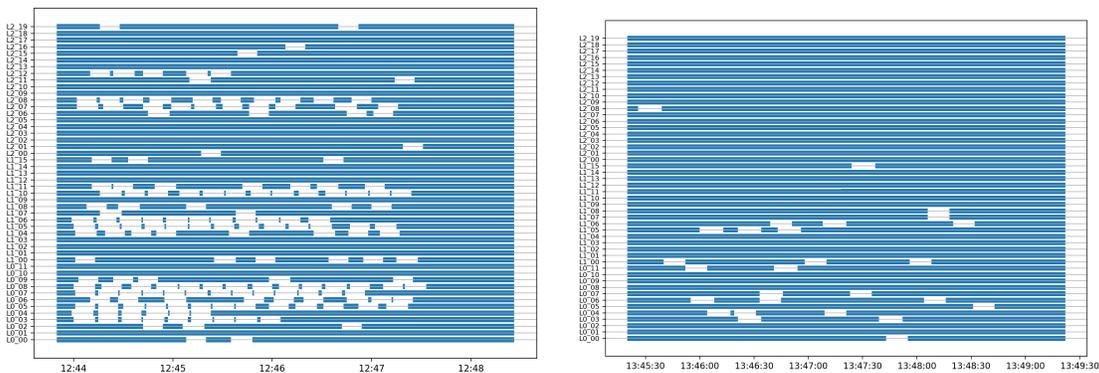


Figure 55: Auto-recovery status of each stave as a function of time for data collected during blue beam only studies. Each horizontal line is a different stave and goes from the innermost to outermost layers with increasing y. A stave in auto-recovery appears as white while a stave fully functioning is shown in blue. The auto-recovery status is shown before the collimators were tightened (left) and after they were tightened (right).

The final test performed with the beam was to increase the value of β^* from 0.7 to 5. This was

performed using both beams and was also found to decrease the rate of auto-recoveries compared to data collected during collisions with the normal β^* value and can be seen in Fig. 56.

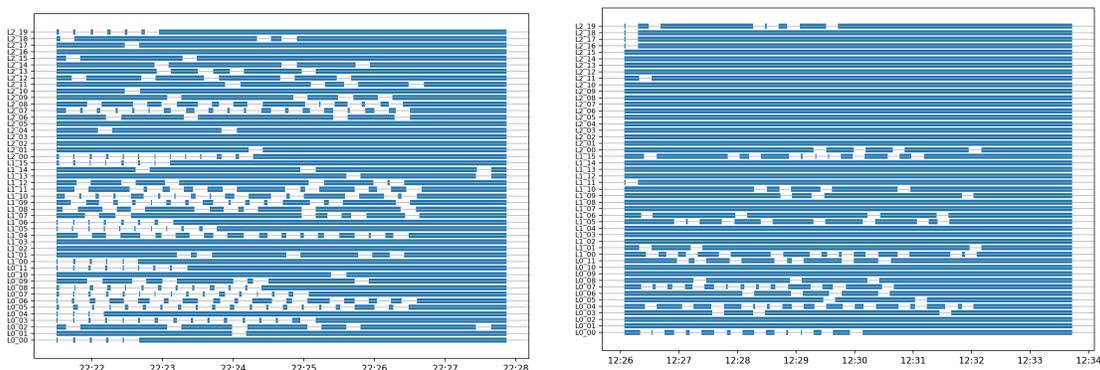


Figure 56: Auto-recovery status of each stave as a function of time for data collected during Au+Au collisions. Each horizontal line is a different stave and goes from the innermost to outermost layers with increasing y . A stave in auto-recovery appears as white while a stave fully functioning is shown in blue. The auto-recovery status is shown during the default value of β^* (left) and the larger value of β^* (right).

9.1.7 Remaining Commissioning Tasks

The MVTX group views the period before physics data taking as occurring in two phases. The first phase involves no beam in RHIC and will be focused on detector operation with cosmic rays. This involves taking a large data sample for internal and also global alignment, if other subsystems are also taking data concurrently.

During this period efforts will be directed towards finalising the detector control software. This software is already in use by the MVTX group but minimal control has been handed over to the shift crews. At this moment they are capable of monitoring the detector health, starting and ending runs, and recovering certain clock errors. The ability to turn the detector on or off, and configuring the MVTX for data taking is performed by the MVTX group as, without changing beam conditions, the detector can be left on for extended periods of time. As a part of the control software, the auto-recovery process will be streamlined by removing unnecessary WAIT commands or ALPIDE reconfigurations. The aim is to reduce a single stave recovery from approximately 12 s to less than 3 s, however this can not be verified until we have beam returned to RHIC.

Due to the unexpected conclusion of Run 23, a final set of operating parameters for the MVTX in terms of configuration and Au+Au beam conditions was not determined. The results presented in 9.1.6 demonstrate the necessity of continuing the study during the next opportunity for beam. Specifically, the collimator test was conducted using a single Au beam so it would be of the highest importance to repeat this during collisions when both Au beams have reached an emittance of 0.56 mm-mrad. From this point it should be possible to determine the emittance value at which the collimators can be tightened to allow the MVTX to operate within sPHENIX. To maximise the physics opportunities with sPHENIX, it would be informative to perform a RHIC fill but avoid bringing the beams into collision until the stochastic cooling has been able to reduce the emittance

to a point where the collimators can be tightened and the MVTX is set to take data. This would then allow for a full eight hour physics fill with all detectors recording events instead of a five or six hour physics fill with all detectors.

It is expected that a combination of improved recovery processes and closer collimators will bring the MVTX to a point where it can take data of the quality expected from sPHENIX. As the beam backgrounds are suspected to originate outside of the sPHENIX hall then they must have significant energy and an increased threshold is expected to have a minimal impact on reducing the background while an increased value of β^* will significantly reduce the luminosity in the IR and so altering both of these values is expected to have a lower impact on the backgrounds observed and thus would be tuned after optimising the first two methods.

9.2 INTT

The silicon strip INTermediate Tracker (INTT) plays key roles in the sPHENIX experiment; 1) it provides seeds and enhances the resolution of the track reconstruction of the charged particles, and 2) it has a unique capability of fast time response of 60 ns allowing to readout collisions data from each single RHIC's beam bunch-crossing and suppress event-pileup background. The INTT detector consists of three main components:

1. The first main component is the active region containing the silicon strip sensors. Four sensors are wire bonded to 52 FPHX chips, which are wire bonded and epoxied to two High-Density Interconnects (HDI), and both are epoxied to a Carbon-Fiber-Composite (CFC) stave for structural support. This assembly is known as a ladder. The ladder covers an acceptance of pseudorapidity of ± 1.1 . The total active area of silicon sensors on one ladder is 91 cm^2 . On each ladder, four thermistors were mounted allowing us to measure temperatures across and monitor the stability of the silicon detector during operation.
2. Several ladders are arranged tangentially in a cylindrical shape around the beam pipe and are referred to as layers. The active area of the sensors is approximately half the width of the HDI, which lends itself to a configuration where two layers are required to obtain hermetic 2π azimuth coverage, and it is called a barrel. The detector consists of four layers located at approximate radial positions of 7.2, 7.7, 9.7, and 10.3 cm from the beam axis. The layers are held in a position with a series of CFC end-rings.
3. The service barrel, attached to the end-rings of the barrels, houses services like cables, and cooling/nitrogen. It consists of CFC tubes and aluminum brackets. The cables connect to the ladders and readout cards (ROCs), and cooling/nitrogen connect to manifolds; both of which are located about 1.5 m from the IP in Z. The service barrel allows the detector to be rolled in one unit. The support structure includes a carbon fiber tube and aluminum I-beams. It facilitates INTT detector installation and operation.

In conclusion, the INTT detector is composed of two half-silicon strip barrels of cylindrical shapes almost 10 feet long. It is mounted along the beam pipe and centered at the collision's interaction point in the sPHENIX experiment.

Before addressing the commission's phases, we would like to point out the following remarks:

- Prior to the detector construction, there were several studies of the prototype BNL silicon strip ladders at the Fermilab Test Beam Facility in 2017 and 2019, and at Tohoku University Test Beam in 2021. During these beam tests, ladder assembly and readout chain components were improved to reach the reliable best performance in the detector response and assembly.
- Prior to the detector installation in the sPHENIX IR, we had several rehearsals of the detector installation using final assembly fixtures, crane, and realistic conditions in the assembly lab at the physics department at BNL.

After the installation of the INTT detector in the sPHENIX experiment on February 28, 2023, the detector went through several commissioning phases which can be summarized as follows:

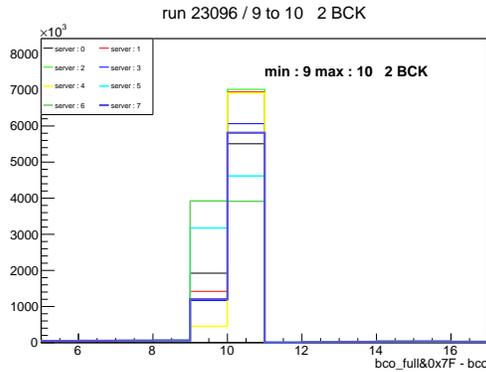
9.2.1 Commissioning of Cooling and Nitrogen Systems

The FPHX chips mounted to the ladder produce heat (64 mW/chip) which needs to be removed from the interaction area. Because of its location, a sub-atmospheric system was developed to provide liquid cooling to the ladders. Each half INTT barrel produces 94 watts of heat and is made up of 28 individual ladders. We remove heat by providing each ladder with 150 ccm of water (plus an additive for corrosion). The calculated flow for each half barrel is 4.2 lpm. To provide this flow, each half barrel utilized its own Chilldyne Cdu 15 sub-Atmospheric circulator to move liquid through the system. Heat is then removed from the liquid by the conventional chiller. This chiller is set to 15 C° which provides 18 C° temperature liquid at the detector located 100 feet away. During commissioning, we were able to establish 5.5 lpm of flow to each of the half barrels. More than adequate to maintain the ± 1 C° and this provides very stable conditions and very low noise level for the INTT detector during operation. Furthermore, to avoid condensation on the surface of silicon sensors, each end of the barrel is supplied with dry nitrogen to keep the surface of the sensors dry and clean. Multiple flow meters supply this nitrogen to ensure even distribution. This keeps the area below a dew-point of 5 C°.

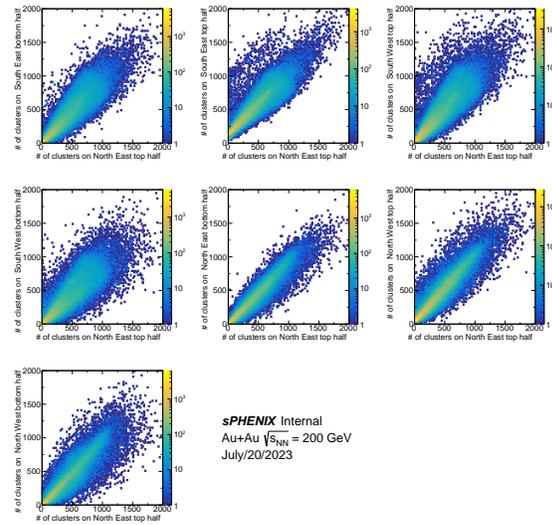
9.2.2 Commissioning of Low Voltage and High Voltage Systems

Commissioning of Low Voltage (LV) and High Voltage (HV) systems is critical to ensure the normal operation of the INTT detector. The latter is a silicon semiconductor detector very sensitive to external noise. Consequently, the LV and HV systems have to include noise filter modules to isolate the INTT detector from external noises, e.g., from rack power and other sub-detector electronics in the experiment. The LV and HV systems were commissioned as follows:

- **Low Voltage System:** The LV distribution comprises low voltage power supplies QPACS modules housed in MegaPAC cases, LV filters, and power transport cables. The MegaPac and QPACS are off-the-shelf products. The unipolar distribution modules and LV filters were developed and made at BNL. One module filter has 40 channels of pi-filters and channel relays. The INTT uses 32 channels from each module, 16 analog power channels, and 16 digital power channels. These modules are powered using digital and analog power QPACS. During the commissioning, this system was adopted to INTT needs by mapping filter outputs to match the INTT detector and making new power cables. The LV filter crates and modules were tested in the Silicon-lab rack and in INTT's racks on the platform to ensure proper functionalities. The voltage on each channel output was measured and checked several times using a multimeter. The continuity of power cables and connector mapping were checked before cable installation. The voltage measurements were performed on the detector end of the power cable.
- **High Voltage System:** The HV system comprises bias power supplies with 16 channel modules housed in Wiener Full crate MPODs; HV filter boxes; and power transport cables. The HV filters are passive pi-type filters with an RC filter cut-off frequency of 1.6 kHz. The Zener diode on the channel input limits the max voltage to 200 Volts. Sixteen boards are housed in one box. All boards were bench tested. After box assembly the connectivity and short circuit tests were performed for all connectors. After filter boxes were mounted in racks and connected to corresponding power supplies, the voltage on each channel output was measured and set to 10 Volts to allow testing with a multi-meter prior to use in the INTT detector.



(a) Internal FPHX chips counter relative to the GTM counter



(b) Clusters correlation data obtained using seven FELIXes vs reference FELIX. Each FELIX collects data in one part of the INTT detector (see text for details)

Figure 57: Measurements obtained in the INTT detector from RHIC Run-23 in minimum-bias Au+Au collisions at $\sqrt{s_{NN}} = 200$ GeV (run number 20444 and magnetic field ON).

9.2.3 Commissioning of FELIX Readout System

The FELIX readout system’s duties are to configure the FPHX chips and the ROCs with the desired parameters; to select and transfer collision data from the ROCs to computer memory; to ensure the whole detector is synchronized with itself and with the beam and to check the health of the detector. The FELIX readout can choose hits of interest in two different ways: triggered and streamed. The triggered readout selects the data from a collision upon receiving a trigger from the Global Timing Module (GTM). The streamed readout transmits all hits coming from the chips without any sort of selection. During the installation of the detector, we tested its condition by taking calibration runs in streamed readout mode. That allowed us to successfully validate both the installation itself and that readout mode. After installation, the detector was timed-in and operated in triggered mode. The process involved multiple iterations of the FELIX firmware and it was ultimately done by scanning the readout’s 4 timing parameters in order to maximize the number of hits accepted per collision, as well as to improve clock alignment between the chips and signal-to-noise ratio. In the streamed readout mode the online timing process is substantially simpler since the complexity of INTT’s event building is transferred from the FELIXes to the offline analysis. Figure 57(a) shows the relative timing between the FPHX chips and the GTM counter. Different colors show different FELIXes, with all eight of them being shown. The better the alignment of all the FELIXes the easier it is to distinguish hits from different RHIC’s collision bunches. After weeks of timing work, all of the chips are now aligned within two RHIC beam crossing periods. Figure 57(b) shows the clusters of all the FELIXes but #5, which was used as a reference. The plots shown in Figure 57 attest that the detector has consistent synchronization with itself and the beam collisions. Both readout modes were successfully exercised and validated by other analyses presented throughout this report.

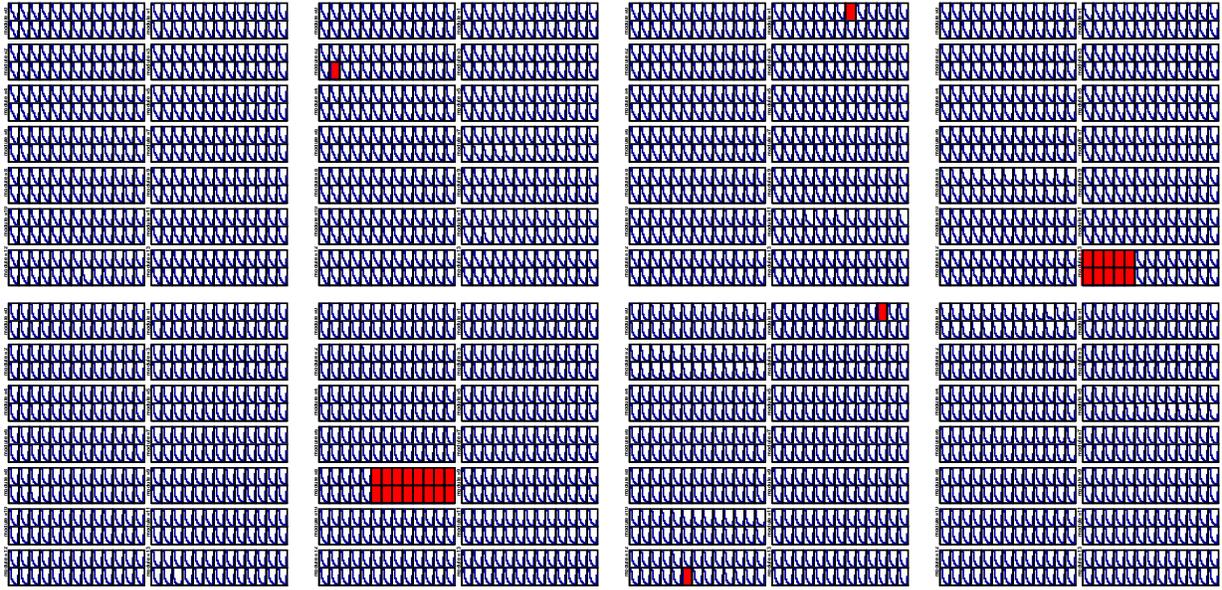


Figure 58: Measured ADC distribution in all INTT's detector chips obtained in Au+Au collisions at $\sqrt{s_{NN}} = 200$ GeV indicating 99% of live channels (using run number 20856).

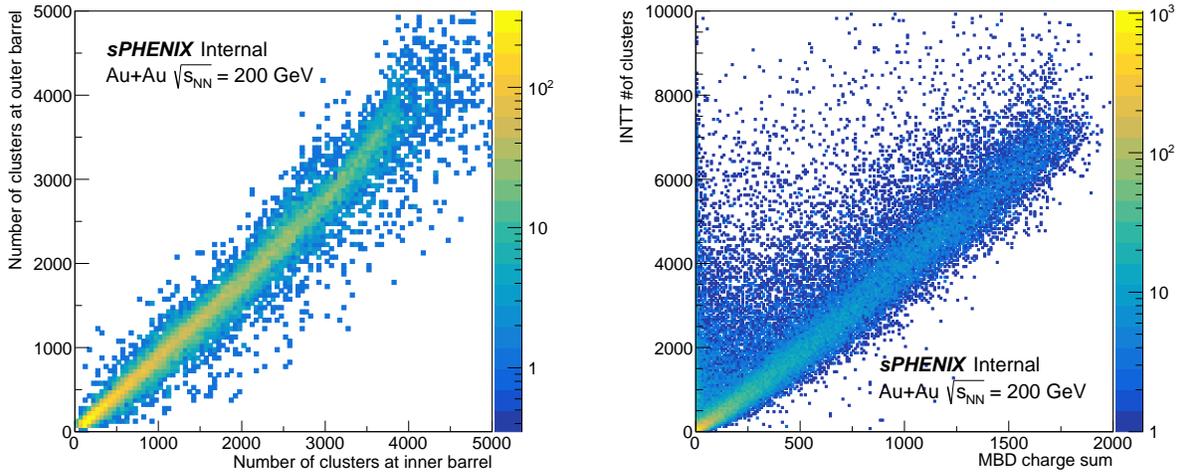
9.2.4 Detector Performance with Beam Collisions Au + Au at 200 GeV

After declaring the "Physics Store" by the RHIC collider with Au+Au at $\sqrt{s_{NN}} = 200$ GeV, the low and high voltages were turned ON for the INTT detector. During that time, the cooling and nitrogen were already ON and operating properly. In the beginning, the main goal of the INTT crew is to time-in the silicon detector relative to the RHIC's clock collisions.

After weeks of timing work, all the FPHX chips are aligned within two RHIC beam crossing periods as shown in the figure 57(a). To prove that the detector was timed-in, it was necessary to observe hits adc distribution in every FPHX chip and charged particle's hits (clusters) in the INTT detector.

Figure 58 shows adc distribution in all FPHX chips of the INTT detector obtained in Au+Au collisions $\sqrt{s_{NN}} = 200$ GeV. These adc distributions reveal that the fraction of active channels remained stable, and it is about 99.0% live channels of 372736 total channels of the INTT detector. The 1% missing channels are represented by the red boxes reflecting the two nonfunctional silicon sensor modules caused by a lack of bias voltage. The bias voltage cables were disconnected at the two silicon module connectors themselves during the installation. The measurement showed that the two silicon modules are healthy and all channels are working properly but elucidate a high noise. Unfortunately, these two bias voltage cables are buried in the barrels, and hard to access two the connectors.

Figure 59(a) shows the excellent agreement of cluster's correlation between inner and outer barrels of the INTT detector obtained using data Au+Au collisions at 200 GeV. This observation indicates that the INTT detector is working properly on its own. Furthermore, this correlation confirms the proper alignment of the packets of information read out from the INTT detector.



(a) Full INTT barrels clusters correlation: inner barrel vs. outer barrel

(b) Clusters in all INTT barrels vs. total charge in the MBD detector

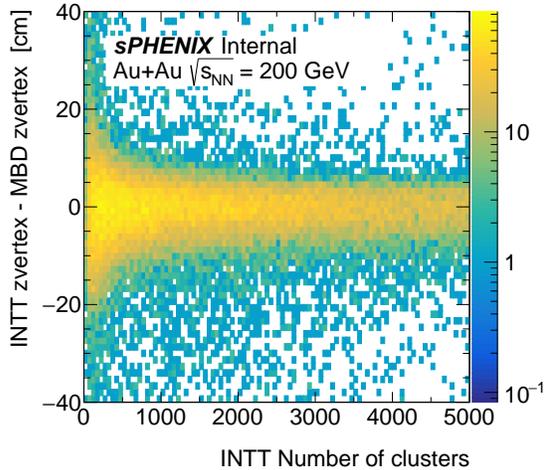
Figure 59: Clusters correlation measurements obtained in the INTT detector from RHIC Run-23 in minimum-bias Au+Au collisions at $\sqrt{s_{NN}} = 200$ GeV and using sPHENIX DAQ big-partition. (a) First full INTT barrels clusters correlations using data run 20444 with magnetic field ON. (b) All clusters in the INTT detector correlated with the total charge in the MBD detector using run number 20708.

Figure 59(b) shows the correlation between hit clusters in the INTT barrels and the total charge in the MBD detector. The correlation again shows the ability to align the two independent detector systems offline using information from their front-end electronics. The INTT detector has an acceptance of pseudorapidity of ± 1.1 and is made of four silicon strip layers this adds capability in the determination of the collision's vertex knowing that the strip size in the ϕ direction is $78 \mu\text{m}$ and 1.6 cm (and 2.0 cm) in the z -direction. Figure 60(a) shows the difference of vertex obtained by the INTT detector and the MBD detector as a function of INTT clusters. The difference is centered at zero indicating the excellent agreement between the two independent detectors in the measured vertex. Figure 60(b) presents the capability of the INTT detector for the determination of the collision's vertex in a high multiplicity environment in the Au+Au collisions using tracks of charged particles observed in the silicon strip barrels. The measured distribution shows clearly that the beam collisions are happening at -15 cm in z -vertex.

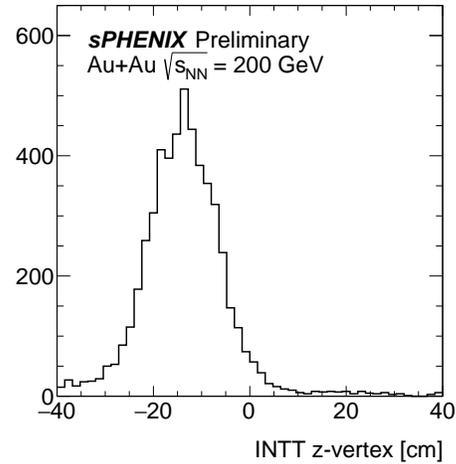
9.2.5 Detector Performance with Cosmic-rays

Since the end of the RHIC Run-23, the sPHENIX sub-detectors were set to take cosmic rays data in the big-partition daq and use the outer HCAL as a cosmic-ray trigger. The principal objective of the INTT cosmic-ray data is to achieve internal alignment of the ladder within the INTT detector itself. Furthermore, the global alignment of the INTT barrels relative to other sub-detectors because the cosmic-ray per event was observed throughout different sPHENIX sub-detectors. Figure 60(c) shows the event-display of one track of cosmic-ray going through all four layers of the INTT detector indicating that the detector is working properly and as expected. This figure was

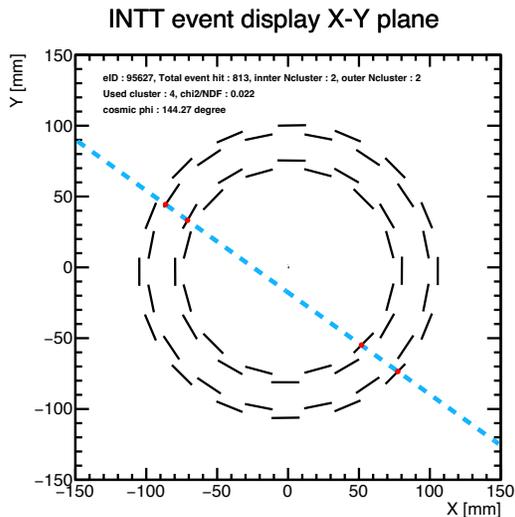
obtained in streaming readout mode (no trigger). Figure60(d) similar to Figure60(c) but the data was obtained in trigger readout mode; the OHCAL detector was used to trigger on cosmic rays.



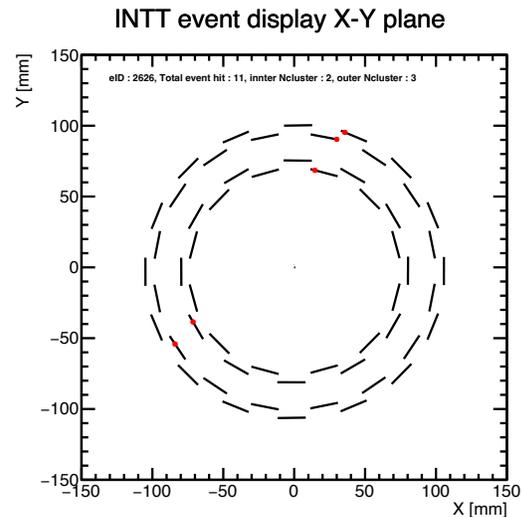
(a) Forward versus backward clusters correlation in measured in the INTT detector



(b) Vertex distribution obtained by the INTT barrels detector



(c) Track of cosmic-ray observed in the four INTT layers obtained in **streaming readout mode**, run number 25566 and magnetic field OFF. The cyan line presents a straight-line fit.



(d) Track of cosmic-ray observed in the four INTT layers obtained in **trigger readout mode**. The OHCAL used as the cosmic-rays trigger, run number 25184 and magnetic field ON.

Figure 60: The INTT detector performance obtained from RHIC Run-23 in minimum-bias Au+Au collisions at $\sqrt{s_{NN}} = 200$ GeV and using sPHENIX DAQ big-partition. (a) Forward and backward clusters correlation using full INTT detector obtained with run number 20444, and magnetic field ON. (b) Collisions vertex distribution was measured using tracks in the INTT barrels detector obtained from data of run number 20708 and magnetic field ON.

Table 2: Summary of the commissioning of the INTT systems

INTT Summary		
System Commissioning	Status	Improvements with Beam Run-24
Cooling and Nitrogen	done ✓	none
Low/High voltages	done ✓	none
FELIX Readout	done ✓ <ul style="list-style-type: none"> • All eight FELIXes stable and validated • Det. aligned 2 beam RHIC bunch crossing 	<ul style="list-style-type: none"> • Validate streamed readout
Detector	done ✓ 99 % live channels	none

9.2.6 Summary

The commissioning of the cooling/nitrogen system, power system, FELIX readout system, and silicon strip detector was achieved successfully. The detector has 99 % live channels. The performance results of the INTT detector obtained with the Au+Au collisions at 200 GeV and cosmic rays were excellent, stable, and above our expectations. The INTT detector is ready for data taking in the RHIC's Run-24, see table 2.

9.3 TPC

9.3.1 TPC Overview

The sPHENIX TPC is a tracking detector with a radial coverage that extends from approximately 20 cm to 80 cm in radius, -1.1 to $+1.1$ in η , and covers the full azimuth in ϕ . The detector's principle components that shall be discussed here include:

- One field cage (inner & outer field cages, a central membrane, and two wagon wheels).
- Twenty four GEM modules in three sizes (R_1 , R_2 , R_3) for a total of seventy two modules.
- A gas system capable of operating in closed loop or pass-through modes.
- A sub-atmospheric chilled water system to remove heat from the electronics.
- Two laser calibration systems (termed the "diffuse" and "line" lasers).
- 624 Front End Electronics or FEE cards supplied with DC LV power and optical readout.

Aspects of the TPC system including DC LV power, Felix Card, Dataflow/storage are common to several sPHENIX detector systems and discussed elsewhere on the document. Here we report the commissioning status of each item in the above list, as well as broad system operating goals. As an up front executive summary these are presented in Table 3.

Task	Status	Comment
Field Cage to 45 kV	Achieved	Routinely tested to 55 kV
Gas System	Achieved	Used throughout the run in pass-through with flows of 5 lpm and 10 lpm.
Cooling System	Achieved	Electronic temp stable w/o air locks.
FEE Live > 95%	Achieved	99.7% live FEE measured.
FEE throughput > $10 \frac{GB}{sec}$	Achieved	Long term stability up to $300 \frac{GB}{sec}$
FEE noise < 4ADU	Achieved	Typically 2.1 ADU
Slow Controls	Achieved	Full monitoring and logging of all TPC-relevant parameters.
GEMs to 4850 Volts	Achieved	The addition of fast trip circuitry allowed operation at full GEM voltage. Nonetheless, the high voltage stability is significantly worse than tests performed prior to transport/installation and must be investigated between runs.
Tracks	Partly	Single event displays below final operating voltage show track evidence. Beam lost before full operating voltage applied.
Diffuse Laser	Partly	Laser induced hits observed, but uniformity of response requires additional investigation.
Line Laser	Not Achieved	Must be work subsequent to the establishment of tracking, which was only partly achieved.

Table 3: Status of TPC Commissioning Goals

Table 4 shows the list of tasks that should be achieved for the TPC prior to the 2024 run. The GEM deformation, GEM operation, Diffuse Laser and Line Laser tasks must be accomplished in the order shown. Gas system and FEE tasks can be accomplished in any order and/or in parallel with the other tasks.

Task	Description
GEM deformation	Indirect evidence that the GEMs foils are deformed was collected. If this turns out to be the case, access to the TPC face is required to address the cause and apply a solution.
GEM operation	Stable operation operation only be achieved after the deformation issue is addressed. This exercise will validate the gas choice, HV configuration, and protection procedures.
Diffuse Laser System	Must understand the variation in response across the central membrane area which may or may not be coupled to the GEM issues.
Line Laser System	Must follow the establishment of stable GEM operation.
Task	Description
Gas in Closed Loop	Demonstrated in principle, but slated for operation only after TPC signals are stable.
FEE run zero-suppressed	Can be developed in parallel with all of the above. Can be developed on the bench, but long term stability required the full suite of on-detector electronics.

Table 4: Tasks required during shutdown. The top part of the table separates out those tasks whose cause is beyond simply running out of time when the beam was lost.

9.3.2 Detector turn-on and commissioning

The TPC was shipped from Stony Brook University on December 20, 2022 and delivered to the Physics building to await installation in sPHENIX. The TPC was delivered to the 1008 complex on Jan 17, 2023, followed by installation and hook up as shown in Figure 61. The central membrane HV and GEM HV systems were checked for continuity via a precision current measurement made at low voltage. The operation for the working gas and full HV was approved on May 18, 2023, which marked the start of the TPC commissioning.

The TPC electronics was installed at the construction site of Stony Brook University in Nov 2022. A health checkout and fixes was performed after the detector delivery to BNL, and after the detector was installed in its final location in sPHENIX. Signal fiber routing and back-end DAQ (FELIX and EBDCs) were operational on April 17, 2023, followed by a link to the upstream GTMs on April 17. Maximum throughput of the readout was demonstrated to be above 300GB/s as shown in Figure 62, though in practice it is limited by the DAQ and file system to around 10 GB/s.

A systematic check out of the electronics healthiness and problem fixing was performed, resolving about 20 different issues. 99.7% of FEEs (622 out of 624) were operational at the completion of the sPHENIX installation.

The electronics have been in operation mode since end of Apr 2023, with most of time under

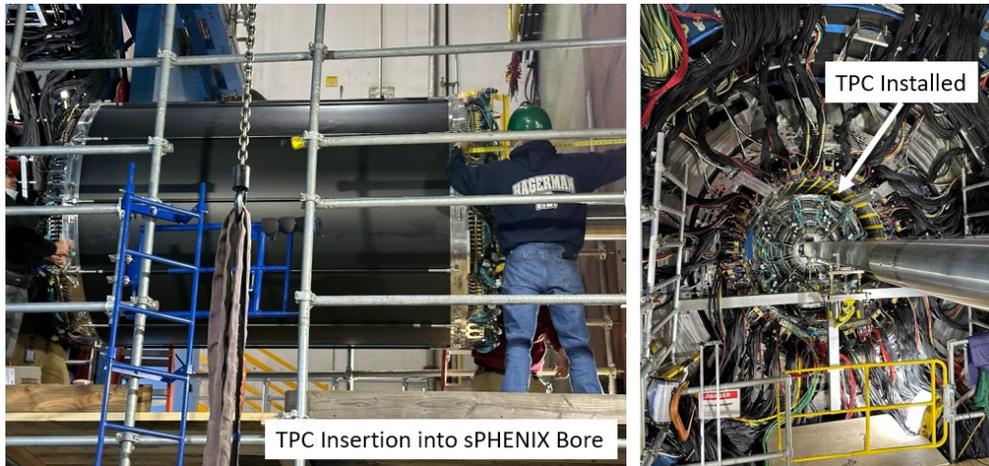


Figure 61: The TPC was inserted into the sPHENIX Bore in January 2023 by rolling it into position while supported by a temporary I-beam spanning the length of the sPHENIX bore.

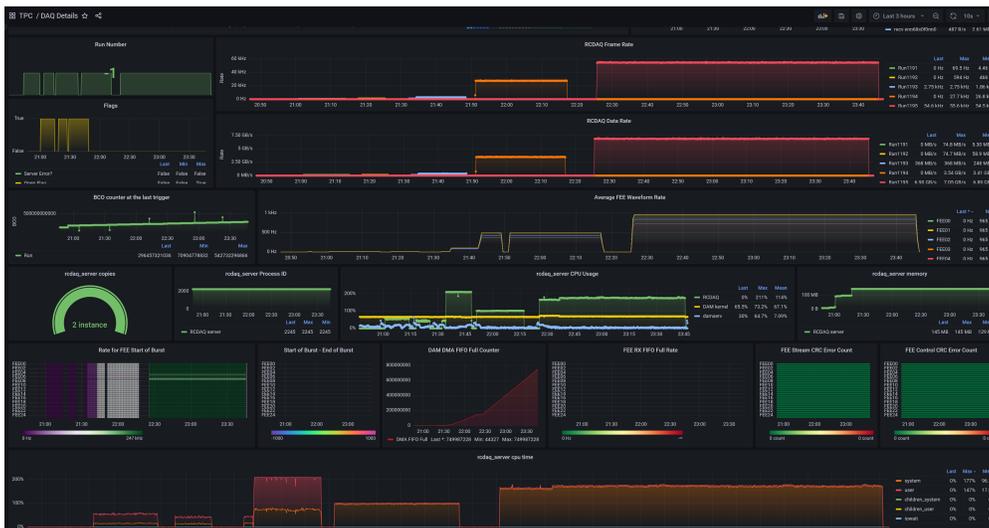


Figure 62: Each stage of the TPC Readout pipeline is monitored in detail with about 20k metrics recorded every 15 seconds to a Time Series database. The analytic and monitoring are provided by sPHENIX Grafana service. This was the pipeline monitoring during the maximum readout throughput test at 340 GB/s (48 PCIe endpoints with 7 GB/s each) non-zero suppressed full detector data readout to server memory.

constant monitoring runs. The electronics noise was found to be very small (2 ADU for most channels) and stable as shown on Figure 63. The fraction of live electronics channels is shown in Figure 64 as a module-by-module map indicating not only excellent live channel counts overall, but also from every module.

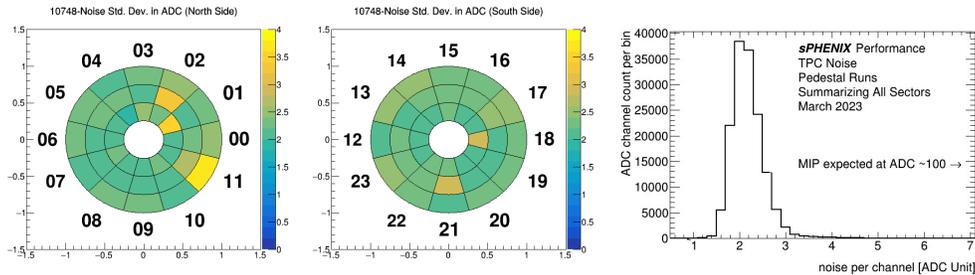


Figure 63: Typical TPC noise measurement, with this particular run taken on June 13, 2023 with GEM at an operating HV (4.875kV).

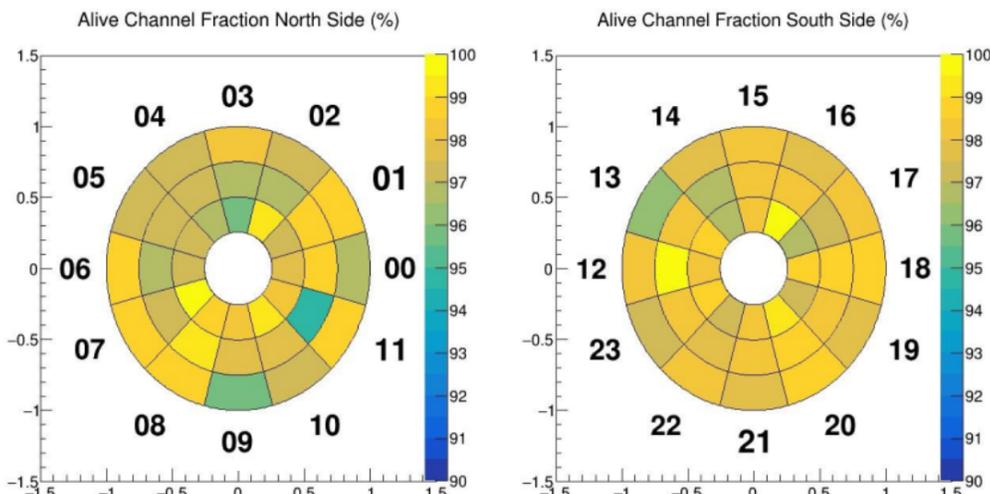


Figure 64: TPC Live Channel snapshot showing not only high live channels overall, but additionally module-by-module.

9.3.3 Commissioning datasets

No-beam data After the initial high voltage turn on without beam, hits are already visible in TPC. Without a cosmic trigger, most interactions are residual radiation in the detector that induced low energy particles traversing through the detector as shown in Figure 65.

First collision data with field off The first collision data, as shown in Figure 66 was taken on June 13 with 100 Hz rate leveling and with a zero crossing angle. The sPHENIX magnet was at a near zero field setting for this run. TPC GEM HV is at 4525V which was likely below the MIP sensitivity. Nonetheless, heavy ion collision produced several highly ionizing tracks which is visible in this display. This data were also validated the calculation of the gas drift speed as shown in Figure 67.

First collision data with field on First collision data with field on followed on June 23 with 100 Hz rate leveling and with a 2-mrad crossing angle, and the sPHENIX magnet at full field. GEM voltage was set to 4.45 kV with a central membrane voltage of 45 kV. Two events from the

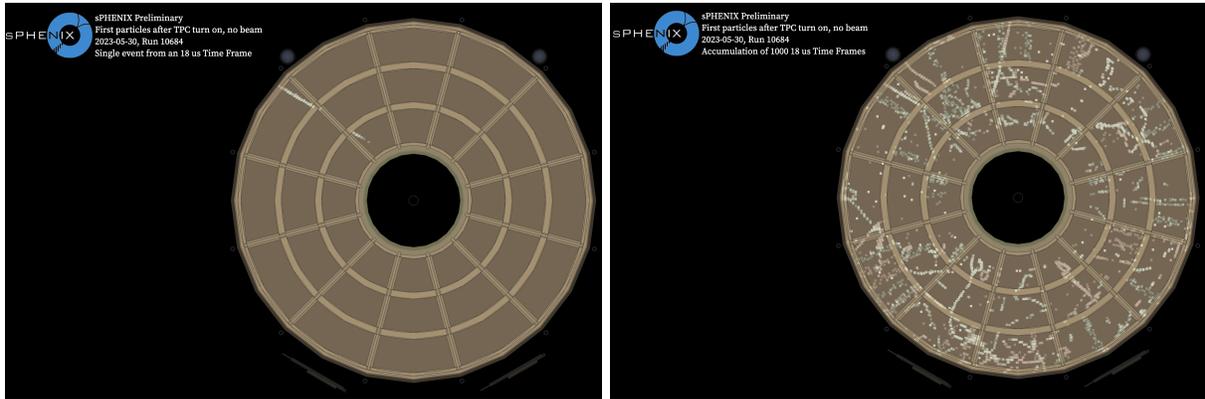


Figure 65: First TPC HV-on data on May 30, for a single TPC Time Frame (one drift time) and overlay of 1000 Time Frames. There was no beam and no magnetic field during this first HV test, and hits which is expected to be dominated by the low energy radiation background.

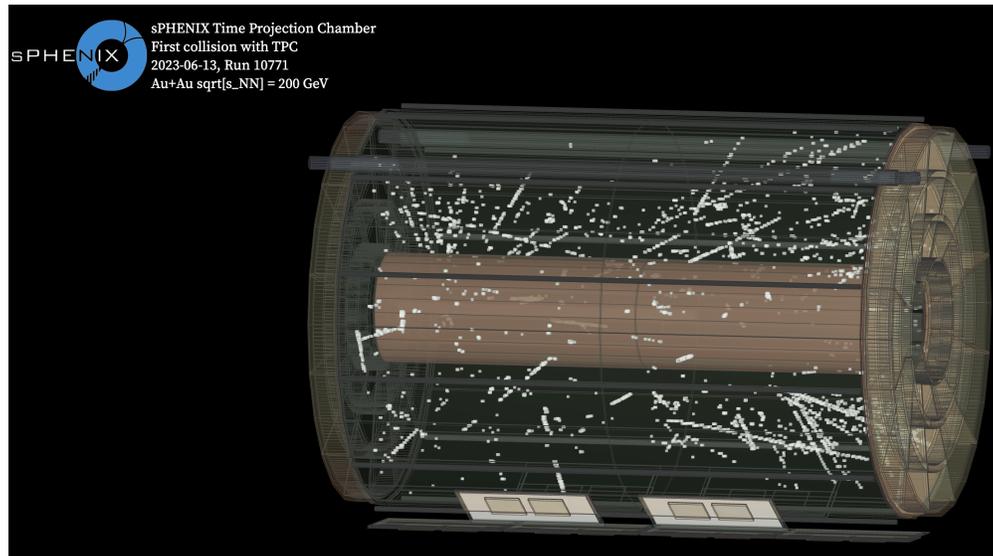


Figure 66: First TPC data with collision. TPC GEM HV is at 4525V which was likely below the MIP sensitivity. Nonetheless, heavy ion collision produced several highly ionizing tracks which is visible in this display. The sPHENIX magnet was at a near zero field setting for this run.

first second of data were displayed in Figure 68.

First diffuse laser dataset Diffuse laser commissioning followed the GEM HV turn on in Mid-June. And five out of the six planned laser heads were turned on and timed in on July 6, 2023. Figure 69, and the full animation at https://indico.bnl.gov/event/20190/contributions/79225/attachments/48895/83238/Animatedhits_laser-3.mp4 illustrate two sheets of electron clouds generated by laser hitting central membrane, drifting towards the end caps for readout. The laser system’s stability and uniformity are still under study with the remaining issues in the next subsection.

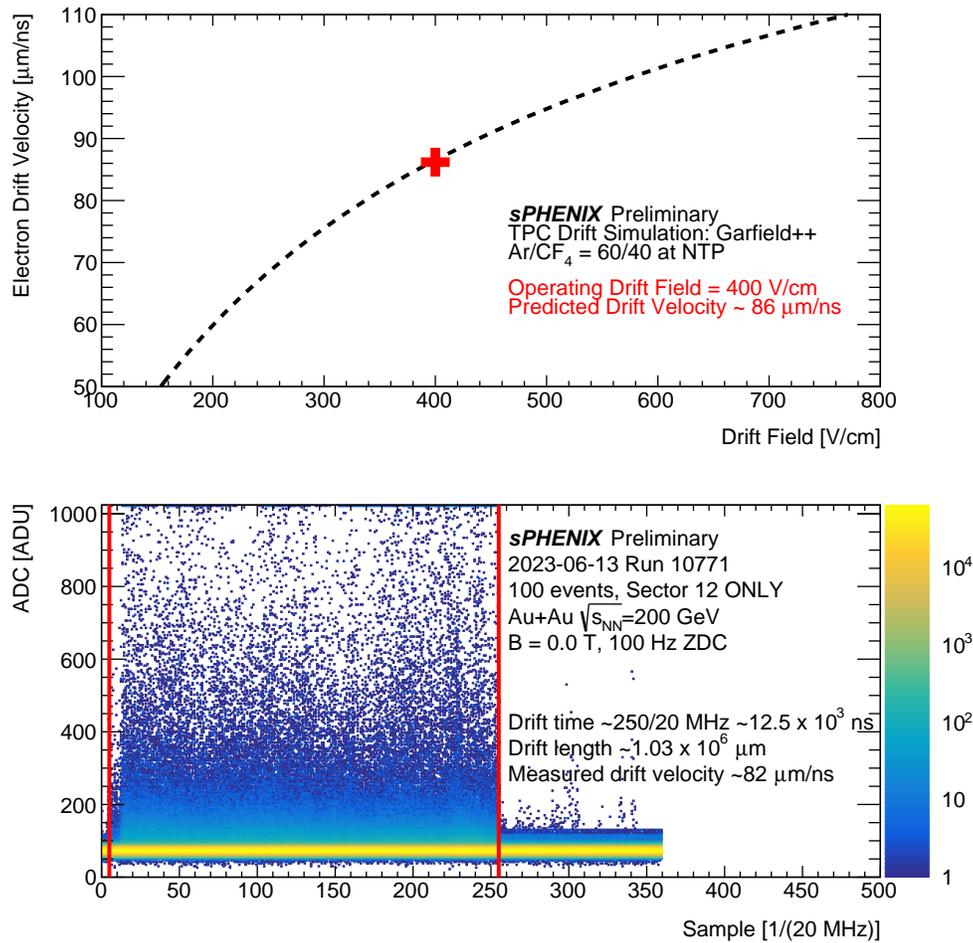


Figure 67: First TPC data with collision were used to validate the calculation of the gas drift speed to within 10 %. Top panel: simulation of electron drift speeds in sPHENIX TPC for different drift fields. The current operating drift field is 400 V/cm - labeled with the red star. Bottom panel: Measurement of electron drift speed in sPHENIX TPC in $\sqrt{s_{NN}} = 200$ GeV Au+Au collisions.

Magnet On Data Following the Beam Loss In the time period following the loss of beam at RHIC, the sPHENIX magnet had field on for a bit more than one week. This period was coincident with the first period of TPC running following the commissioning of the fast trip system for minimizing spark damage (see below). Although the GEMs shows current draw in time with the laser and the spark system picks up the occasional hits from naturally occurring heavily ionizing particles, no signals were found to be present in the readout stream. Debugging efforts to isolate the problem (which the TPC group believes is likely unrelated to the HV instability issues discussed below) were unsuccessful at curing the issues before the cold magnet period ended.

Magnet Off Data Following the Beam Loss After the magnet ramp-down/warm-up and several small TPC repairs, the observation of laser cosmic pulses was restored. Initial operation used 2/3 lasers on the south end lasers and 3/3 lasers on the north end. Although the lasers were not yet

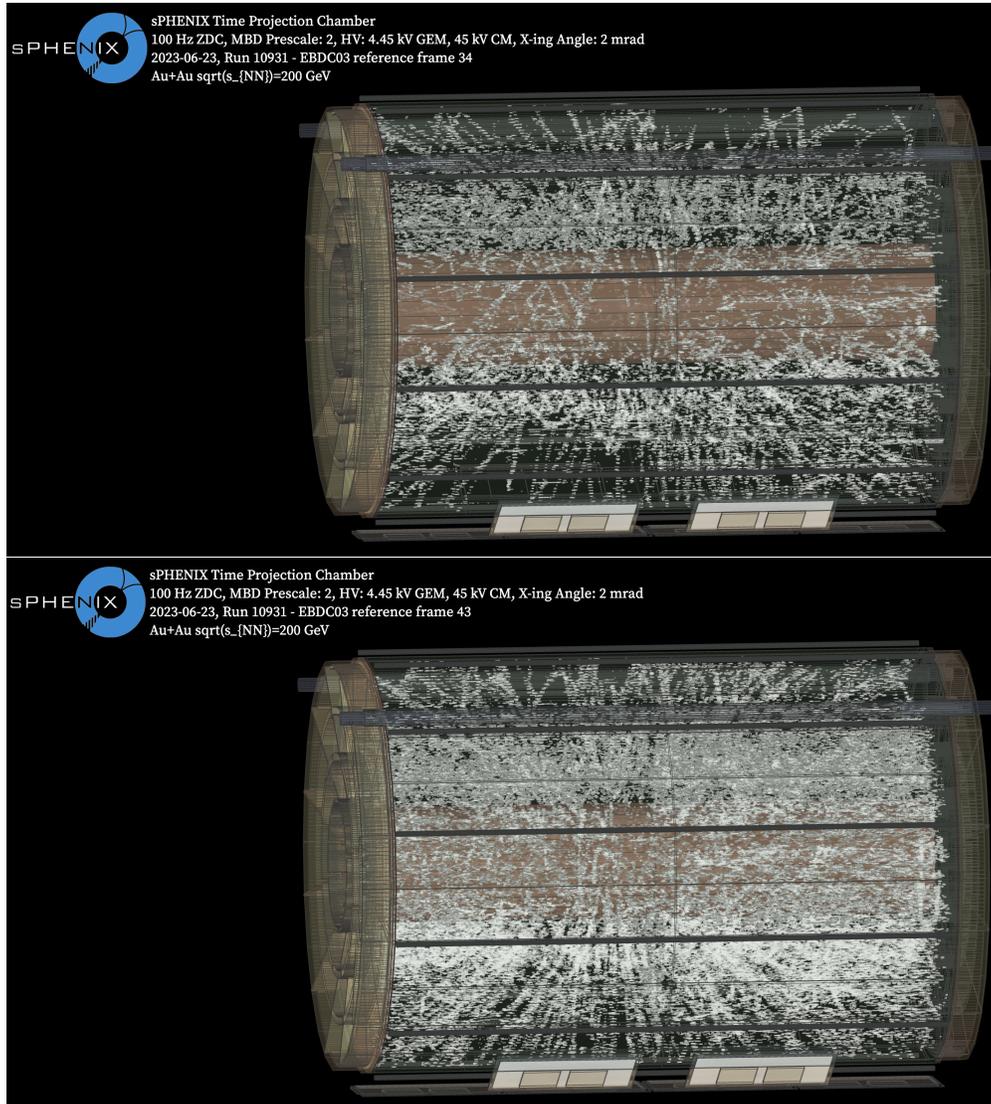


Figure 68: First TPC data with collision with full sPHENIX operating magnetic field. GEM voltage was set to 4.45 kV with a central membrane voltage of 45 kV.

normalized in output light intensity, in-time laser induced pulses were observed in illuminated sectors (top panel of Figure 70) and hit patterns from a few minutes of HCAL cosmic ray triggers (bottom panels of Figure 70).

9.3.4 High Voltage Issues and GEM Deformation

Although the TPC was installed in January 2023 and fully cabled shortly thereafter, commissioning nonetheless began in earnest only in late May following safety approval to both begin the flow of operating gas and to apply high voltage. At this point in time, no access to the TPC face was possible because multiple layers of other sPHENIX systems were installed and fully blocked access to both end caps. Two modules were immediately noticed with a severe issue as drawing excess

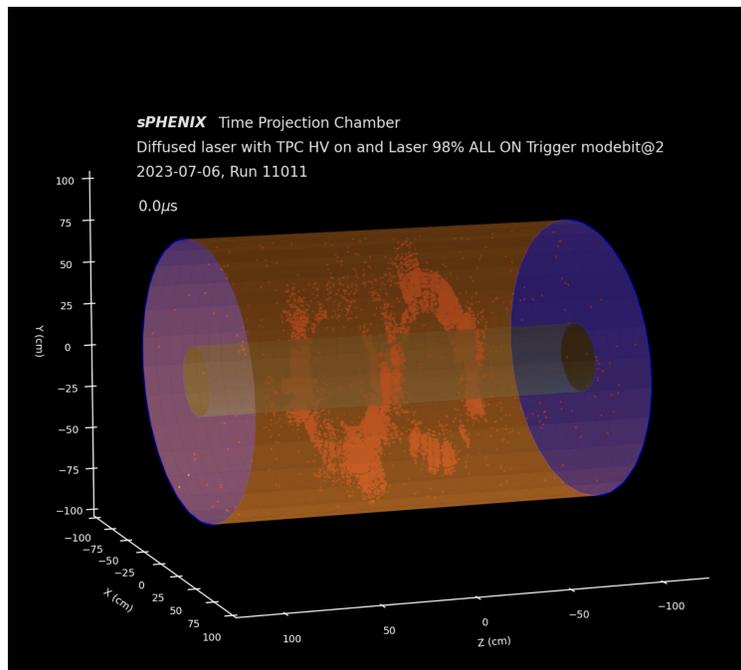


Figure 69: Snapshot of the animation for TPC drift for a single event from the diffuse laser test in run 11011. The full animation can be viewed at https://indico.bnl.gov/event/20190/contributions/79225/attachments/48895/83238/Animatedhits_laser-3.mp4

current. Multiple diagnostic steps were taken to verify that the cause of both these issues related to deformation of the bottom GEM allowing it to touch the readout pads. These included:

- The measured value matches to better than 1% the current draw that would result from a GEM-pad short.
- The short disappears after 30-60 minutes following the removal of power from the FEE (causing a small thermal swelling of the pad plane to subside).
- The shorts drive high noise into a subset of pads whose locale is consistent with a mechanically distorted or wrinkled GEM.
- Pulsing the bottom GEM measures the capacitance of this GEM to the pads and effectively the gap size and verifies the changed distance.
- A meeting was held with the ALICE TPC experts (whose detector is very similar to the sPHENIX TPC) and they uniformly agreed that some mechanically-driven deformation of the GEMs is present inside the TPC now that was not present prior to shipping and installation into sPHENIX.

Under the circumstance of a non-planar GEM, a variety of deleterious effects are guaranteed. Among these are:

- **Voltage instability.** The low IBF mode of the sPHENIX TPC necessitates large electric fields in the transfer gaps. While demonstrated safe for proper GEMs over many hours and independently for every one of the installed modules, a variation of the gap size will certainly generate a high voltage instability.

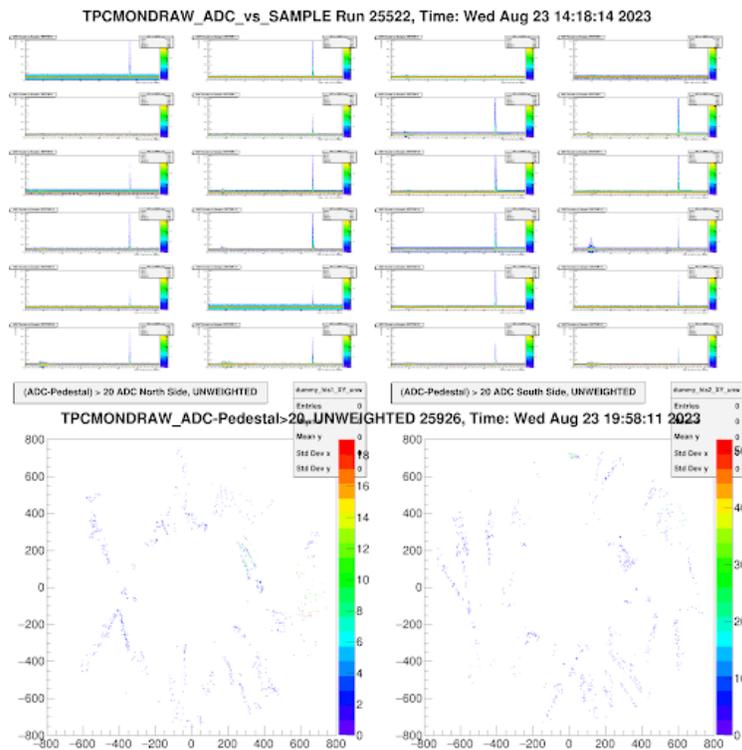


Figure 70: The top panel shows timed laser firing in all illuminated sectors and the bottom panel shows the hit pattern after a few minutes of cosmic ray triggers.

- **Permanent GEM damage.** The sPHENIX GEMs were made with particularly highly segmented top planes to effectively eliminate the possibility of damage from a spark within a GEM. This was demonstrated by thousands of sparks applied once per second to sPHENIX GEMs in a fully energized GEM-stack and in operating gas wherein damage never occurred. Different from a GEM spark, a gap spark increases the voltage in the GEM and opens the possibility of damage.
- **Non-uniform IBF performance.** The high transfer gap fields are a critical aspect of the low IBF operating mode. In the presence of non-uniform gap fields, non-uniform IBF is assured dramatically complicating the precision of the tracking system. The impact of this can only be measured with beam data.

In the face of this hardship and inability to address the root issue, the TPC group in consultation with management embarked on our best path forward:

- Pursue every avenue to operate the TPC under its current condition.
- Attempt to minimize the spark damage to a few percent of the active area while nonetheless bringing most of the modules to operating voltage. Presently, the permanent damage sits at just below 1.2% of the active area ($\frac{18 \text{ dead}}{1536 \text{ stripes}}$).
- Pursue engineering investigations of the source of the deformation and the between-runs mitigation.

We implemented and commissioned spark detection systems to identify sparks in the face of high standing currents and set modules to standby voltages. The penultimate incarnation is capable of digitizing the full waveform of a spark via capacitive pick-off of high speed waveforms from the resistor chain. This produced “stable enough” operation to achieve full voltage (but not nearly so stable as measured before delivery).

Engineering investigations pinpoint no single cause of the trouble, but nonetheless suggest simple solutions to alleviate several points of mechanical over-constraint and moreover to relieve the modules by simply loosening and re-tightening the bolts to each module. None of these solutions is possible until between beam periods where we can get access to the face of the TPC.

9.3.5 Remaining Commissioning Tasks

- Investigate the cause and solution to potential GEM deformation in the installed configuration.
- Validate the working point for the GEM stack, including gas choices, HV configuration, and ramp/park/spark-protection procedures.
- Establish zero-suppressed data operation.
- Turn-on the line laser system and produce static distortion calibration data.

9.4 TPOT

9.4.1 TPOT components

The Time Projection chamber Outer Tracker (TPOT) is a tracking detector located inside the EMCAL and below the TPC. TPOT aims to include an additional reference point for the particle's trajectory on the outside of the TPC. This additional reference point will greatly improve the ability to calibrate the TPC in conjunction with the reference trajectories provided by the rest of the sPHENIX Tracking System, in particular the correction of the beam-induced space-charge distortions.



Figure 71: TPOT detector as installed on the inside of the EMCAL on December 12 2022.

The detector is made up of 8 modules where each module consists of two Micromegas chambers stacked, one to measure the coordinates along the longitudinal direction (z) and the second along the azimuthal direction (ϕ). Each module is equipped with its own FEE boards (one per chamber) and cooling system. The modules are arranged to provide partial coverage of the bottom three (out of 12) TPC sectors as shown in Fig. 71. The bottom four modules are arranged to cover the full longitudinal stretch of the TPC. Additionally, there are two modules on each side, extending to East and West of the bottom most center sector. These cover approximately half of the longitudinal extent of the TPC. In total, TPOT represents approximately 8% of the TPC acceptance.

Each Micromegas chamber consists of a drift electrode, a micro-mesh, a resistive layer and a readout layer arranged in parallel. The drift electrode is operated at negative voltage. The micro-mesh is located 3mm away from the drift electrode and connected to the electrical ground. The resistive layer is located $128\mu\text{m}$ away from the micro-mesh and is operated at a positive voltage. The readout layer is located immediately below the resistive layer and it consists of 256 parallel strips per chamber. The strips are connected to the front-end electronics and read the signal collected on the resistive layer via capacitive coupling. The chambers are circulated with a 95/5 mixture of Ar/iC₄H₁₀. Each chamber is operated at a gas flow rate ranging from 50 to 100 ccm.

The Front End Electronics (FEEs) of TPOT are mounted directly on the modules and readout using optical fibers. There is one FEE board corresponding to each chamber for a total of 16 boards. Each uses 8 32-channels SAMPA ASIC chips to amplify, shape and digitize the analog signal collected on the strips. There are three low voltage lines per FEE with two of them corresponding to the digital and one corresponding to the analog part. The FEE boards used for TPOT are identical to those used for the TPC.

The FEE are cooled using a subatmospheric-pressure cold water circuit. The water is circulated through aluminium cooling plates mounted on top of each FEE board and separated by a thermally conductive gap pad. A Chillyne unit is used to circulate the subatmospheric pressure water. TPOT and the TPC share the same unit. Water is provided to each module independently and monitored by its own flowmeter. The cooling plates of the two FEE boards of a given module are mounted in series. The system relies on subatmospheric pressure to ensure no liquid leaks outward.

Lastly, the fire safety system of TPOT consists of tubes connected to the middle of each sector. Those allow for both the detection of build up of flammable gas at the bottom of the EMCAL as well as an effective path to inject Nitrogen to displace the built up flammable gas.

TPOT was installed in sPHENIX on December 12, and fully cabled by end of January 2023. The authorization to operate the detector with its working gas and full HV was given on May 18, 2023, which marked the start of the TPOT commissioning.

9.4.2 Commissioning of TPOT High Voltage

The High Voltage system of TPOT consists of one negative drift voltage and four positive resistive voltages for each chamber. This makes a total 80 HV channels.

For the safety of the chambers, tests were done before and after installation. Each CAEN High Voltage channel and High Voltage cable was tested individually to ensure their reliability and absence of leakage current, by operating them through a voltage divider at a voltage significantly larger than that needed for detector operation, and for an extended period of time. No significant issues were found. With the confirmation of the reliability of the system, each channel has been connected to the chambers and tested again, this time to validate the ability of the chambers to hold high voltage. The drift channels were tested for up to 500V for stability and the resistive channels were tested up to 430V, making sure no channel drew prohibitively high leakage current. These tests were conducted repeatedly to ensure consistency before assembly, after assembly and after installation. All but 2 (out of 64) resistive channels could be ramped up to 430V. 2 channels could not be operated above 360V and 380V respectively due to discharges. No problem has been found with the 16 drift channels.

To protect the chambers from possible discharge-induced damage, TPOT is equipped with a spark protection system through the CAEN High Voltage unit. If a given channel draws more current than an adjustable limit (typically a few μA , channel-dependent), for more than a given amount of time (also adjustable and of typically a few seconds), the voltage is ramped down, and the channel is turned off. The shift crew has been provided with a GUI that allows to recover tripped channels manually with a single click, and without the need for contacting a detector expert.

Currently, TPOT operates with three main modes: OFF, SAFE and ON. When in the OFF mode, the channels all hold 0V. When in SAFE mode, the channels hold low enough voltage that there is no amplification in the chambers. SAFE mode is typically used when there is no stable beam in the accelerator. When in ON mode, the chambers are all brought up to operating voltage to collect data from collisions with optimal gain and stability. The shift crew is provided with a GUI to switch between SAFE and ON modes at will. The operation lasts less than a few minutes. The exact values of HV, trip current limits, ramp and more are invisible to the shift crew and are decided on by the experts. The detector HV have been operating routinely since May 20. HV operation has been in the hand of the shift crew since May 25.

After establishing the SAFE and ON operating voltages in the end of May, the frequency of trips was in the order of one or two each day, exclusively when the detector was ON and with beam. The majority of these trips are believed to be due to collision-induced highly ionizing interactions in the detector volume rather than detector instability. No trips were observed in SAFE mode or in the absence of beam. Since Aug 8th, there have been no trips and TPOT has been continuously taking cosmic data with no issues.

9.4.3 Commissioning of the Front-End electronics cooling

A total of 5 temperature sensors are readout for each FEE boards: 4 on the board's PCB, and one on the FPGA. For safe operations, measured temperatures should not exceed 80°C for the FPGA and 60°C for the PCB. With the sub-atmospheric pressure cooling water provided by the Chillyne unit at a flow rate of about 100ccm, TPOT FEE's have been operated continuously since early May 2023, with temperatures below 40°C for the FPGA and 35°C for the PCBs, for most of the time. This is well below the above limits. For the two occasions during which the temperatures of some boards were found to slowly increase over time (around July 4 and July 23), careful re-balancing the flow rates between the 8 TPOT parallel lines fixed the issue.

The FEE cooling system and FEE LV power supply are interlocked so that the FEE are turned OFF if the cooling system is not working properly. This is also part of the TPC setup. It was verified that the interlock works properly at several occasions during the run.

9.4.4 Commissioning of Front-End electronics and Data Acquisition

TPOT uses the same Front-End boards as the TPC. There is one board per TPOT chamber, reading its 256 channels using 8 32-channels SAMPA chips. Commissioning the FEE for TPOT includes:

- deciding the SAMPA chip operation parameters in terms of gain, shaping time, sampling frequency and number of samples;
- quantifying noise levels and ensuring that they remain stable over time
- quantifying at which trigger rates the FEE can be operated in a stable manner

Concerning the SAMPA chip operation parameters, it was decided to use the lowest possible gain (20mV/fC) to avoid saturating the pre-amplifiers too often and the longest available shaping time (160ns), to match the typical signal width in the Micromegas chambers. The sampling frequency

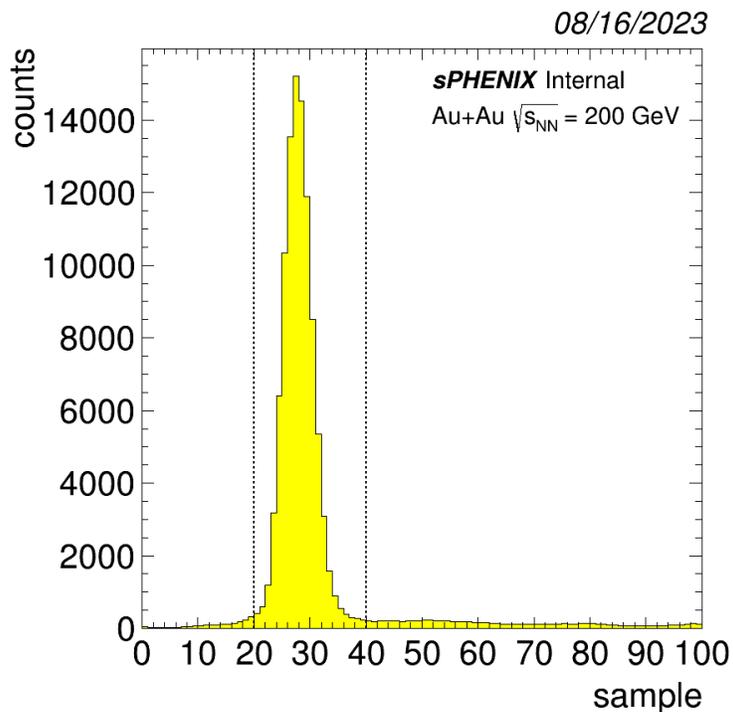


Figure 72: Signal counts above threshold as a function of time sample in one TPOT FEE board. The vertical lines correspond to the physical signal from particle crossing the detector, in time with the trigger. One sample is 50ns.

is 20MHz (50ns per sample). With these settings the TPOT electronic signal is well contained within 20 samples ($1\mu\text{s}$), as illustrated in Fig 72. We took data using a larger number of samples (either 50 or 100 samples per trigger), to monitor the proper timing of the signal with respect to the trigger as well as possible baseline shifts.

Concerning noise levels, the measured RMS is around 8 ADC units, across all channels, and stable over time. This is larger (by a factor 2 to 4) than that measured with the TPC, because of the significantly larger detector capacitance (150pF to 200pF as opposed to 10pF for the TPC GEM detectors). Measured noise levels are consistent with those measured in the lab in a clean environment, and those reported by the SAMPA manufacturer for such input capacitance.

When recording 50 samples by trigger, the TPOT FEE could be operated in a stable manner at a trigger rate of up to 5kHz, limited at that time by another of the sPHENIX subsystems included in the data taking. Another factor of two can be achieved by reducing the number of recorded samples down to 25, or even 20, provided that the signal is properly timed with respect to the trigger. This would bring TPOT stable operation trigger rate close to sPHENIX goal of 15kHz.

So far TPOT FEE and back-end have been operated in triggered mode only. In order to make the maximum of the coming pp run, TPOT FEE and backend should also be operated in streaming-readout mode, along with the other tracking subsystems. This has not been tested yet.

9.4.5 TPOT working point

In absence of magnetic field, all TPOT modules could be operated at the same voltages as that defined in the lab, and for which 100% detection efficiency was measured: -100V for the drift electrode and 400V for the resistive layers. This to the exception of 2 (out of 64) resistive layer channels, which could not be brought to voltages larger than 360 and 380V respectively, above which they would draw currents significantly larger than the allowed trip limits. This corresponds to about 3% of the TPOT acceptance. It must be noted however that even at these low operating voltages, the detection efficiency is not zero, and so the loss of acceptance is not complete.

In TPOT, the sPHENIX longitudinal magnetic field is perpendicular to the drift direction of the primary electrons deposited by particles in the Micromegas drift space. The magnetic field consequently curves the trajectory of the primary electrons in the transverse plane, resulting in a non-zero value for what is known as the *Lorentz* angle to gas-detector experts. Too large values of the *Lorentz* angle result in deteriorated performance for the detector both in terms of spatial resolution and efficiency. For a given magnetic field, the Lorentz angle is reduced by increasing the drift electric field. This was achieved by increasing the drift voltage from -100 first to -300V then -400V, after realizing that -300V was likely not large enough.

A first estimate of the in-situ detection efficiency of the Micromegas detector is obtained by using the correlation between the number of clusters measured in each detector of the same module. In absence of magnetic field and for a drift voltage of -100V, full efficiency is reached at a voltage of about 380V for the resistive layer. This is consistent with more precise measurements performed in the lab. In presence of magnetic field on the other hand, because of the increased drift voltage (here -400V), and the subsequent reduction of the micro-mesh transparency to primary electrons, full efficiency is reached starting from 390-400V for the z detectors and 400-410V for the ϕ detectors. At 360V, an efficiency ranging between 40 and 50% is measured, varying from detector to detector. At 380V, it ranges between 60 and 70 %. These efficiency estimates have an uncertainty of about 10 % due to the poor control over the reference tracks provided by looking at the other view only. More accurate estimates will require correlating the TPOT measurements to that of other tracking subsystems and redo the resistive layer voltage scans.

To summarize the current TPOT acceptance: out of the 64 resistive layer channels (4 per chamber), two channels (one half of a chamber) are completely inefficient, because not properly connected to the FEE. This happened during installation, and could not be fixed once the detector is in place; one channel operates at about 40 to 50% efficiency, and another at about 70% efficiency. A third channel was also operated at 380V (70 % efficiency) for a while because it was developing some small dark current over time, but so far we have had no problem operating it at nominal voltage (400V) either. All other channels are active and fully efficient. This results in an effective acceptance for TPOT of about 96%.

9.4.6 Correlation between subsystems

The left panel of Fig. 73 shows the correlation between the total number of clusters measured per triggered Au-Au collision in the TPOT z detectors, and that measured in the ϕ detectors. A good correlation is expected between the two since the acceptance of the two sets of Micromegas detectors overlaps almost entirely. The correlation is linear. The spread around the diagonal is

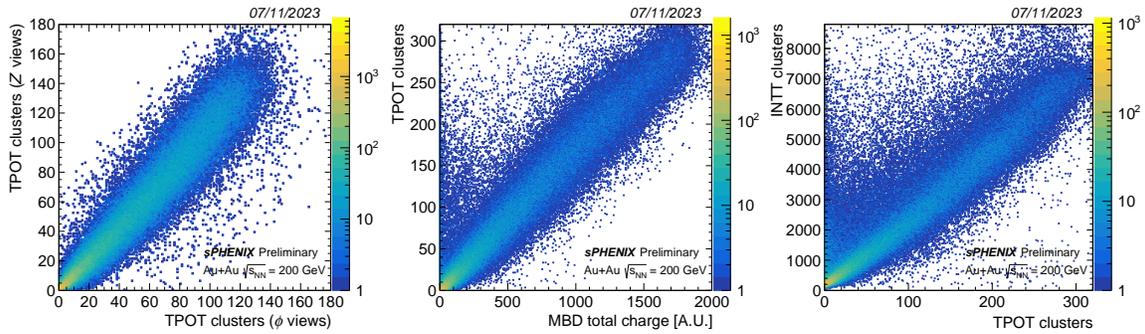


Figure 73: Left: internal correlation between the number of clusters in the z and ϕ TPOT detectors; Center: correlation between the total number of clusters measured in TPOT and the charge (A.U.) measured in the MBD; Right: correlation between the number of clusters measured in the INTT and that measured in TPOT

due to slight difference in active areas, the detectors not being fully efficient, imperfections of the clustering, occasional noise hits. In addition, the slope is above unity because of the two ϕ detectors that are operated at a lower voltage, as mentioned in the previous section.

The center panel of Fig. 73 shows the correlation between the total number of clusters measured in TPOT and the total charge (arbitrary units) measured in the MBD. There is no overlap between the TPOT and MBD acceptance. The correlation between the two subsystem is entirely due to the fact that they measure the same collisions in a synchronized manner, each with its own centrality.

The right panel of Fig. 73 shows the correlation between the total number of clusters measured in the INTT and that measured in TPOT. This time the observed correlation is due in part to the common event activity and in part to overlapping active areas.

These correlation plots demonstrate our ability to combined the data collected by the corresponding subsystems in synchronized full events corresponding each to the same triggered Au-Au collision. They also prove that the majority of the hits recorded in the corresponding subsystem corresponds to actual particles rather than uncorrelated noise.

9.4.7 Summary

Tab. 5 summarizes the status of the commissioning of the TPOT components and what remains to be done before the next data taking.

10 Tracking

The offline track reconstruction has been developed over the last several years in preparation for data taking [10]. To meet physics and CPU requirements, the software was rewritten leveraging common tools available in the wider HEP/NP community. For example, tools developed by the ALICE collaboration and the A Common Tracking Software (ACTS) toolkit have been integrated into the sPHENIX software stack. The track reconstruction workflow, shown in Fig. 74, has

Component commissioning	Status	Remarks
TPOT HV/LV	done	none
FEE cooling	done	none
FEE and DAQ	nearly done	Need to further reduce the number of samples and increase trigger rates; test streaming-readout operation
Detector working point	done	none
Stable operations	nearly done	Online monitoring and diagnostic tools are in place. Missing precise standing orders to the shift crew
Offline processing	in progress	Needs robust event combining and synchronization with other subsystems

Table 5: Summary of the commissioning of the TPOT systems

been continuously exercised over the last 12 months as a part of the simulation Mock Data Challenge. The momentum and distance of closest approach resolution of the track reconstruction in simulation are shown in Fig. 75 as determined for primary charged pions embedded into 50 kHz minimum bias Au+Au backgrounds. The requirements to achieve the sPHENIX upsilon and heavy flavor physics program, which is most dependent on tracking performance, are a DCA resolution of $\mathcal{O}(10)\mu\text{m}$ for charged particles with $p_T > 1 \text{ GeV}/c$ and a momentum resolution of less than $\sim 2\%$ at $p_T \sim 10 \text{ GeV}/c$. Both of these performance goals are achieved in simulation with the current tracking reconstruction.

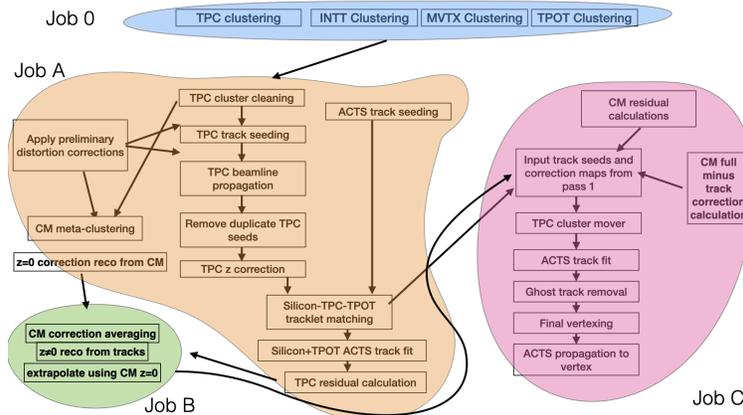


Figure 74: The full track reconstruction workflow, starting from clustering hits to the production of the final reconstructed physics objects, is shown.

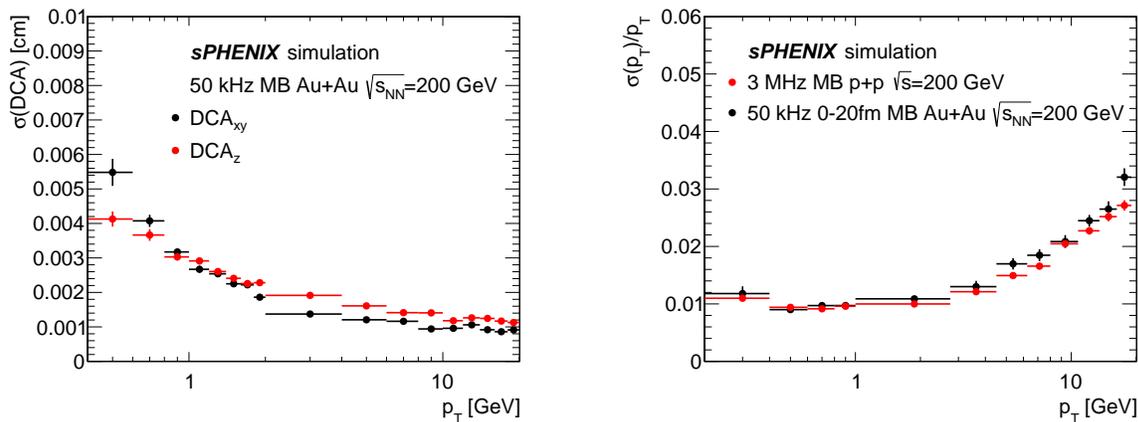


Figure 75: The tracking distance of closest approach (left) and momentum (right) resolution of the offline track reconstruction software in simulated minimum bias Au+Au collisions with 50 kHz pile up.

10.1 Commissioning Track Reconstruction

Due to the unexpected shut down of RHIC early in 2023, the offline track reconstruction was not able to be fully exercised in beam conditions. However, commissioning the online-to-offline reconstruction will continue during beam off time with cosmic tracks. Progress is already being made with single cosmic track events being identified through the MVTX, INTT, and TPOT. We expect only a few weeks of cosmic data taking are needed with all four of the tracking detectors read out together to commission the online to offline workflows and to begin initial detector alignment analysis. Prior to data taking in 2024 we expect to be able to readout all detectors into the offline framework, allowing the processing of several weeks of cosmic tracks. Figure 76 shows an example of an event display showing a track measured using the HCal cosmic trigger, going through the MVTX, INTT, and TPOT. This illustrates the ability of the three subsystems to measure tracks in a clean environment.

11 Offline Computing and Analysis

11.1 Online Monitoring

The online monitoring is based on the PHENIX model which uses a client-server implementation. Each one of the SEBs/EBDCs (see 6 runs a server process which fills subsystem defined histograms with relevant information extracted from events provided by the DAQ. The number of processed events is governed by the speed of the subsystem software — in general the number of events analyzed by the monitoring is small compared to the total number of events recorded. Those servers are designed to run continuously — interventions by the shift crew are not part of normal operations. If a new run starts, the existing histograms from the previous run are automatically saved in files which are shipped to SDCC where a cron job creates web pages for later Q/A and long term archiving. If need be — the servers can be started/stopped centrally by the

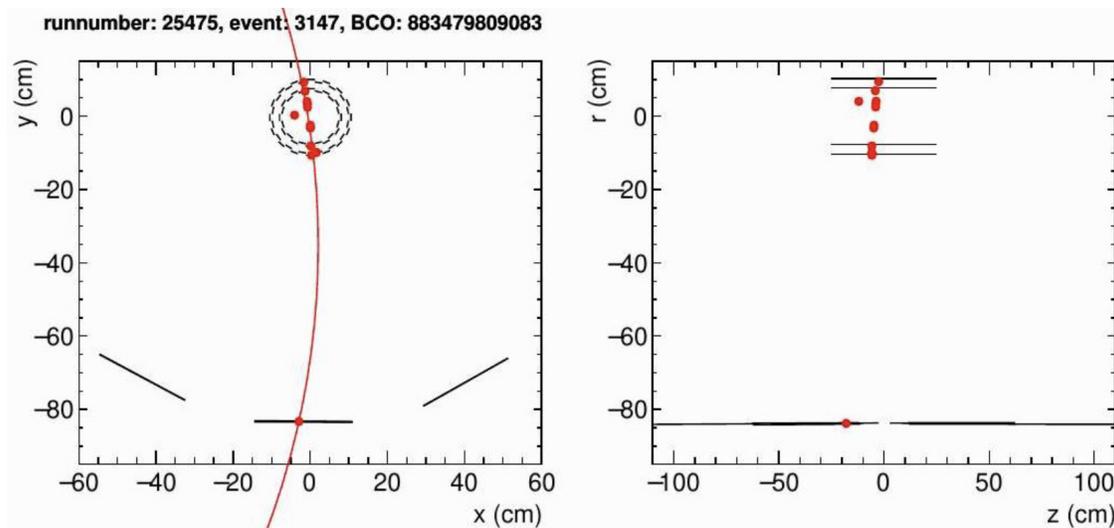


Figure 76: A cosmic track correlating the MVTX, INTT, and TPOT subsystems.

shift crew using a GUI 77 which also indicates if servers have crashed and need to be restarted. Most subsystems are spread over multiple servers (the TPC uses 24 servers) which collect events independently. Checking direct correlations between servers is therefore not meaningful but longer term averages can be used for crosschecks.

The client keeps a list of which servers are associated with which subsystem. To display a monitoring plot, an up to date set of histograms are fetched from the servers of the subsystem. The client then combines the information from each server and displays the result 78. This system is now also used to quickly get a complete picture of the current detector performance for quick tests since it doesn't involve analyzing multiple raw data files with detector subevents. The online monitoring framework has been shown to be stable and new subsystem code is easy to integrate. Most subsystem have at least some basic monitoring in place.

11.2 Data Processing Workflow

Figure 79 shows a simplified flowchart of the sPHENIX data reconstruction workflow. sPHENIX does not employ high level triggers which would require full events, and assembling events is done offline during data reconstruction. We use this to split the event reconstruction into three separate passes where each pass only reads the data it actually processes. Our global event characterization (reaction plane, centrality and low resolution vertex) uses data from the MBD, sEPD, and ZDC. The calorimeter reconstruction handles only calorimeter data from the CEMC and the Inner/Outer Hcals while our track reconstruction processes data from the MVTX, INTT, TPC and TPOT. The track reconstruction itself employs two passes over that data. In the first pass hits which belong to tracks are identified and tracks and residuals are being stored. These residuals are then analyzed and a distortion correction is derived. This distortion correction is then applied in a second pass where the previously identified tracks are refit to achieve the required momentum resolution. The only information which is needed by all passes is our global level 1 information which provides the information needed to synchronize the triggered data with



Figure 77: Online monitoring server GUI

the streaming readout. Every output file contains information which allows us to combine the events from any of these passes during the analysis stage which again enables us to read only the data needed for a given analysis.

The initial version of event assembly software was developed using simulated calorimeter type raw data which showed the feasibility of this new approach (deferring the event building step to the offline reconstruction). Using the data taken during commissioning the software was improved to deal with real data and its problems. In rare occurrences a subsystem drops an event which then requires re-synchronization and leads to loss of events during this step. The current software can handle most of the runs which were taken during commissioning, work is being done to reduce the number of discarded events to a minimum as well as identifying and correcting more failure modes.

The assembly software for the streaming readout as well as combining the streaming data with triggered data is still under development.

This rather complex workflow is managed by a local PanDA [11] instance which is currently being

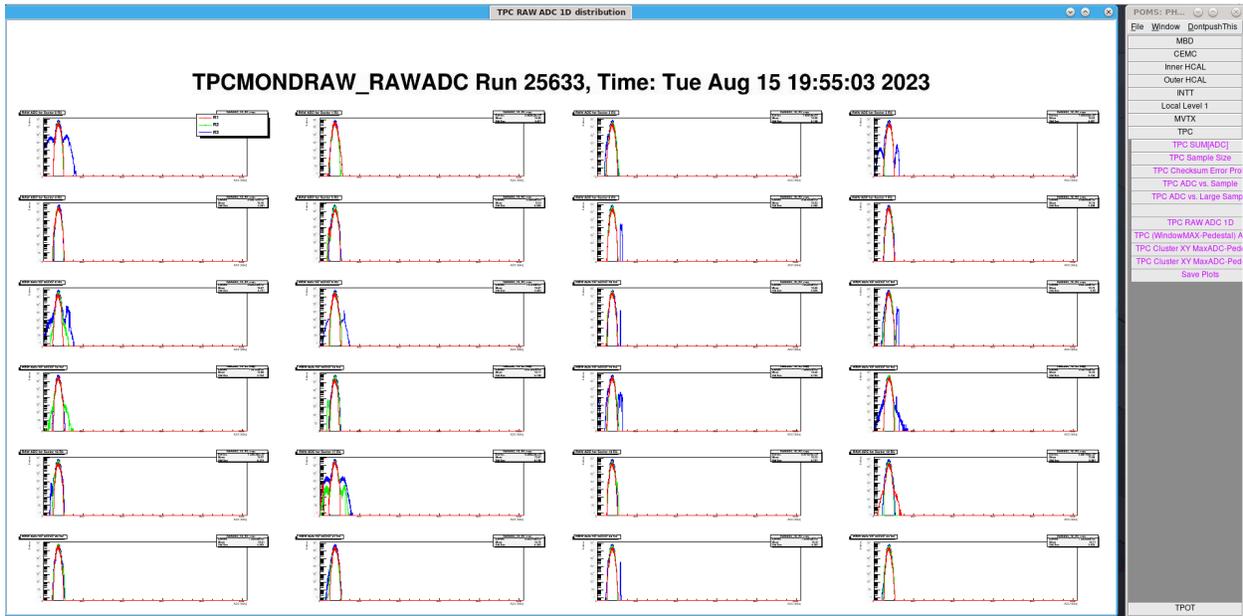


Figure 78: TPC online monitoring plot, displaying results collected from its 24 servers

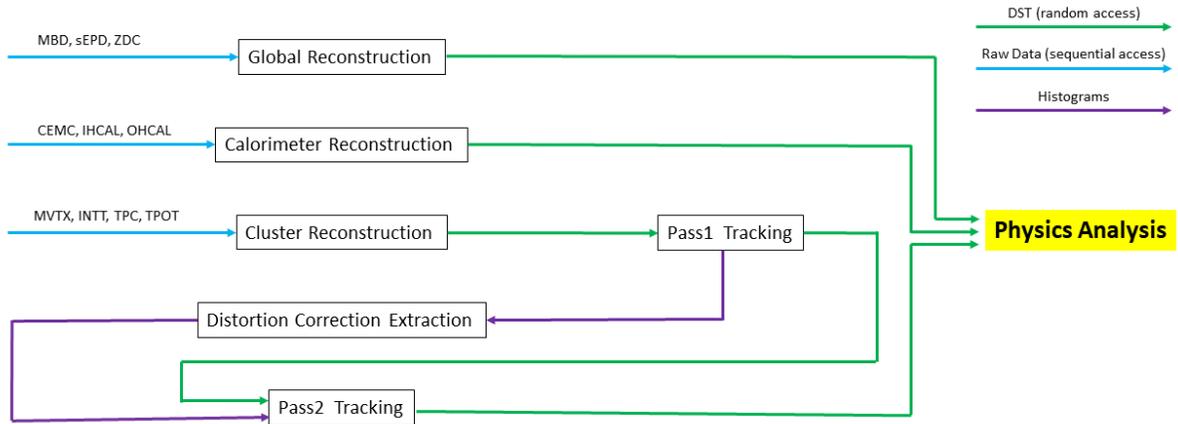


Figure 79: Simplified sPHENIX data reconstruction workflow

commissioned. The calorimeter and global reconstruction pass each of which needs two stages (event assembly and reconstruction) and is combined at the current time has been shown to be functional. We foresee frequent reprocessing of the data sPHENIX took in the last weeks which will soon be managed by PanDA which in turn will give us valuable experience running it.

A Remaining Commissioning Tasks

Many of the sPHENIX commissioning goals were accomplished during Run 2023 at RHIC, but some remain to be accomplished between runs and some can only be completed at the beginning of Run 2024 with collisions. Additionally, information for polarimetry will need to be tracked. The 2024 cold QCD measurements in p+p collisions require for each bunch crossing the online logging of information for proton-spin monitoring and the proper normalization of transverse spin asymmetries. The ZDC and SMD will be used for the local sPHENIX polarimetry, offering a fast diagnostic of the spin state for each bunch crossing. The ZDC and MBD will be used for the determination of the relative luminosity. To that end, the digital scaler information of all three detectors will have to be integrated into GL1. An overview of the tasks is shown in Table 6.

System	Before Run 2024	During Run 2024
Trigger	Firmware and software development of calorimeter triggers	First operation of calorimeter triggers with beam
TPC	<ul style="list-style-type: none"> • FEE firmware completion • tests of zero suppression • completion of MJACK to mitigate SEU • development of digital current • cosmic ray data with and without zero suppression 	<ul style="list-style-type: none"> • Stable operation with HV • collision data with and without zero suppression • testing of digital current and SEU mitigation
DAQ	<ul style="list-style-type: none"> • Tests with zero suppression in calorimeters and TPC • Throughput and livetime tests with multievent buffering • Development of offline event building • Any additional development needed to achieve routine 15 kHz • Improvements in reliability, data integrity, and error handling 	<ul style="list-style-type: none"> • Tuning of zero suppression • Timing of detectors to new triggers • Spin: integrate ZDC, SMD and MBD digital scaler information into GL₁
MVTX	<ul style="list-style-type: none"> • Field off cosmic data for tracking development and alignment • Development of mitigation strategies for background and lock-up 	<ul style="list-style-type: none"> • Field off and field on collision data for tracking development and alignment • Tests of mitigation strategies for background and lock-up
INTT	Field off cosmic data for tracking development and alignment	Field off and field on collision data for tracking development and alignment
EMCal HCal	(HCal) tower-by-tower cosmics analysis	Demonstration of design energy resolution and response uniformity

Table 6: Overview of major commissioning tasks needed with and without beam.

References

- [1] sPHENIX Collaboration. sPHENIX Beam Use Proposal, 2022. URL: https://indico.bnl.gov/event/15148/attachments/40846/68568/sPHENIX_Beam_Use_Proposal_2022.pdf. 1
- [2] Martin L. Purschke. RCDAQ, a lightweight yet powerful data acquisition system, 2012. URL: http://www.phenix.bnl.gov/~purschke/rcdaq_doc.pdf. 6.1.1
- [3] J. Anderson, K. Bauer, A. Borga, H. Boterenbrood, H. Chen, K. Chen, G. Drake, M. Dönszelmann, D. Francis, D. Guest, B. Gorini, M. Joos, F. Lanni, G. Lehmann Miotto, L. Levinson, J. Narevicius, W. Panduro Vazquez, A. Roich, S. Ryu, F. Schreuder, J. Schumacher, W. Vandelli, J. Vermeulen, D. Whiteson, W. Wu, and J. Zhang. FELIX: A PCIe based high-throughput approach for interfacing front-end and trigger electronics in the ATLAS upgrade framework. *Journal of Instrumentation*, 11(12):C12023–C12023, dec 2016. URL: <https://doi.org/10.1088/1748-0221/11/12/c12023>, doi:10.1088/1748-0221/11/12/c12023. 6.1.2
- [4] Kai Chen, Hucheng Chen, Jin Huang, Francesco Lanni, Shaochun Tang, and Weihao Wu. A Generic High Bandwidth Data Acquisition Card for Physics Experiments. *IEEE Trans. Instrum. Measur.*, 69(7):4569–4577, 2019. doi:10.1109/TIM.2019.2947972. 6.1.2, 11
- [5] sPHENIX Collaboration. sPHENIX Technical Design Report, 5 2019. URL: <https://indico.bnl.gov/event/7081/>. 6.2
- [6] C. A. Aidala, V. Bailey, S. Beckman, R. Belmont, C. Biggs, J. Blackburn, S. Boose, M. Chiu, M. Connors, E. Desmond, A. Franz, J. S. Haggerty, X. He, M. M. Higdson, J. Huang, K. Kauder, E. Kistenev, J. LaBounty, J. G. Lajoie, M. Lenz, W. Lenz, S. Li, V. R. Loggins, E. J. Mannel, T. Majoros, M. P. McCumber, J. L. Nagle, M. Phipps, C. Pinkenburg, S. Polizzo, C. Pontieri, M. L. Purschke, J. Putschke, M. Sarsour, T. Rinn, R. Ruggiero, A. Sen, A. M. Sickles, M. J. Skoby, J. Smiga, P. Sobel, P. W. Stankus, S. Stoll, A. Sukhanov, E. Thorsland, F. Toldo, R. S. Towell, B. Ujvari, S. Vazquez-Carson, and C. L. Woody. Design and beam test results for the sphenix electromagnetic and hadronic calorimeter prototypes. *IEEE Transactions on Nuclear Science*, 65(12):2901–2919, 2018. doi:10.1109/TNS.2018.2879047. 8.2
- [7] Megan E. Connors. Test beam results and status of the sphenix calorimeter system. In *2017 IEEE Nuclear Science Symposium and Medical Imaging Conference (NSS/MIC)*, pages 1–4, 2017. doi:10.1109/NSSMIC.2017.8532770. 8.2
- [8] M. Mager. ALPIDE, the Monolithic Active Pixel Sensor for the ALICE ITS upgrade. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 824:434–438, 2016. Frontier Detectors for Frontier Physics: Proceedings of the 13th Pisa Meeting on Advanced Detectors. URL: <https://www.sciencedirect.com/science/article/pii/S0168900215011122>, doi:<https://doi.org/10.1016/j.nima.2015.09.057>. 9.1
- [9] P. Moreira et al. The GBT Project. In *Topical Workshop on Electronics for Particle Physics*. CERN, 2009. doi:10.5170/CERN-2009-006.342. 9.1
- [10] Joseph D. Osborn, Anthony D. Frawley, Jin Huang, Sookhyun Lee, Hugo Pereira Da Costa, Michael Peters, Christopher Pinkenburg, Christof Roland, and Haiwang Yu. Implementation of ACTS into sPHENIX Track Reconstruction. *Comput. Softw. Big Sci.*, 5(1):23, 2021. arXiv:2103.06703, doi:10.1007/s41781-021-00068-w. 10
- [11] PanDA development team. <https://panda-wms.readthedocs.io/en/latest/>, 2020. 11.2